Deconstructing Job Search Behavior*

Stefano Banfi
Ministry of Energy, Chile

Sekyu Choi
University of Bristol

Benjamín Villena-Roldán
CEA, DII, University of Chile

October 25, 2018

Abstract

We use an unusually rich data from a Chilean job board to document various theoretically relevant facts regarding job search. We show how application behavior is influenced by (1) demographics such as gender, age, and marital status, (2) alignment between applicant wage expectations and wage offers, (3) applicant fit into ad requirements such as education, experience, job location and occupation (4) timing variables, including unemployment duration, job tenure (for on-the-job searchers) and business cycle conditions. Our paper provides novel evidence that can discipline current and future search-theoretical frameworks.

Keywords: Online job search, Applications, Search frictions, Unemployment, On-the-job search, Networks. *JEL Codes:* E24, J40, J64

*Email: stefano.banfi@gmail.com, sekyu.choi@bristol.ac.uk and bvillena@dii.uchile.cl. We thank Guido Menzio, Brenda Samaniego de la Parra, Jan Eeckhout, Shouyong Shi, Yongsung Chang, Francois Langot, Bariş Kaymak, and seminar participants at Cardiff University, University of Manchester, Diego Portales University, the 2016 Midwest Macro Meetings, the 2016 LACEA-LAMES Workshop in Macroeconomic, the Search & Matching workshop at the University of Chile, 2017 SAM Meetings in Barcelona, Spain, 2017 SECHI Meeting, 2017 SM3 Workshop, and Universidad de Los Andes Bogotá, Université de Montréal, Universidad de Chile, and the IDSC of IZA Workshop: Matching Workers and Jobs Online for insightful discussion and comments. All errors are ours.

Introduction

What kind of jobs workers look for and how much effort they exert are critical pieces of information for labor market outcomes. These search activities determine wages and job allocations in the economy to a large extent. Key magnitudes for policy debates such as the unemployment rate, mobility, and income inequality are affected by jobseekers' decisions. We economists have some evidence about *extensive* and *intensive margins* of search effort[1], but we know very little about the kind of jobs workers search for, what we call a *selective margin*. As applications are tentative allocations, they are of prime importance to understand ex post outcomes. Hence, while we keep an eye on the search effort, our main contribution is to carefully document nuanced application decisions workers make according to their own characteristics, job requirements, and context, and discuss their implications.

Data from online job boards offer us the opportunity to learn about job search more deeply than probably ever before. While online job search is different from traditional methods, this modality deserves increasing attention as its importance and efficiency has increased over time (Kuhn and Mansour, 2014). In most theoretical models, job seeking behavior is highly stylized: workers accept or reject offers using simple optimal rules. However, jobs are complex objects with many interlinked aspects which matter for job seekers: wages, fit, location, timing, etc. On the other hand, job seekers' characteristics may also affect optimal search strategies.

We use information from `www.trabajando.com`, a job posting website with presence in most of Latin America, in addition to Spain and Portugal. We exploit a comprehensive dataset on daily applications of job seekers to job postings in the Chilean labor market during the period 2008 to 2016. Our dataset contains some novel features. We observe detailed information on both sides of the market: education, occupations,[2] and experience for individuals and for job postings (as requirements stipulated by firms). Moreover, we observe detailed information of seeker and job characteristics, as well as both desired and current wages for individuals (wages of last full time jobs if unemployed) and the wages firms expect to pay at jobs they are posting.

The main research question in this paper is how individuals, facing a set of online job ads, choose to apply to some jobs and forgo others. Given our unique setup, we can deconstruct behavior into two dimensions. First, an intensive dimension, where a particular characteristic

[1] We follow conventional definitions. The extensive margin shows if workers search at all, i.e. if they are classified as unemployed. The intensive margin captures how much searching effort workers exert providing they do search.

[2] We observe a one-digit classification, created by the website administrators.

of the applicant or the posting affects the probability of application, which relates to standard notions in the literature (Faberman and Kudlyak, forthcoming; Gomme and Lkhagvasuren, 2015; Leyva, 2018; Mukoyama, Patterson, and Şahin, 2018). The second dimension is a selective one, where a particular coincidence between characteristics of the worker and the job impacts the probability of application. We disentangle these two aspects by estimating an application decision equation.

We also provide a methodological contribution. To overcome the fact that we only observe effective applications and not the entire set of relevant job positions for each candidate,[3] we use the bipartite *network* formed by job seekers linked through common job applications. Using this network, we construct the choice set of an applicant as the list of all ads applied by seekers linked with her, and create weights given a similarity metric between two job seekers. Our methodology overcomes the problem of only observing actual choices in the data, a recurrent problem to study consumption choices in industrial organization and marketing (Van Nierop, Bronnenberg, Paap, Wedel, and Franses, 2010; Abaluck and Adams, 2017), taking a very different approach.

Instead of this network approach, we could have defined choice sets as segments defined by an arbitrary set of job characteristics as Şahin, Song, Topa, and Violante (2014) and Herz and van Rens (2015), or by clustering algorithms as Banfi and Villena-Roldán (forthcoming). Our procedure has two important advantages. First, it relies on applicant revealed preferences over a probably large number of observable and unobservable (to the econometrician) job characteristics to define similarities among jobs, instead of often few arbitrary dimensions such as locations, occupations, or industry. Using arbitrary choice sets, one would inevitably ignore attempted mobility across segments, a likely important issue as shown in Carrillo-Tudela and Visschers (2014). Second, the network approach allows us to define individual choice sets, generating key variability for parameter identification.

Our empirical exercise reveals interesting patterns with respect to job seekers' application decision making and search effort, measured by the probability of applying to a position in the relevant set. We find that some demographic characteristics are quite relevant at the intensive margin: men apply to more job ads than women, while married individuals apply more than their single counterparts if they are unemployed. We also document that search effort decreases with age, which is consistent with evidence in Choi, Janiak, and Villena-

[3]Not observing page views in the website is a potential shortcoming in the literature using online job board data. Unfortunately, `www.trabajando.com` does not keep records of page views by applicant for two reasons: (i) it is very expensive to keep these records while the information is of little use for the job board operators, and (ii) applicants need to be logged in when viewing job ads, a requirement that would reduce the likelihood of getting applicants into the board.

Roldán (2015) and Menzio, Telyukova, and Visschers (2016), among others. This decrease in effort is of crucial importance for unemployment insurance design, as studied in Michelacci and Ruffo (2015).

One of our main findings relates to how job seekers align themselves with heterogeneous types of jobs. We find that search behavior is highly sensitive to the requirements of educational level, experience, location and occupation requirements and that job seekers target an optimal or most preferred type of job, which is not necessarily the one that matches perfectly their current characteristics: the probability of an application peaks when the applicant is slightly underqualified in terms of education but the pattern is reversed in the case of experience requirements. An implication of this result, is that job seekers' incentives to apply to jobs that are further away from the most preferred job (in terms of these characteristics) decrease, although this decrease is not symmetrical. We also study how this fit evolves over the life-cycle, the duration in the current labor force status and across different levels of aggregate unemployment rates (business cycle conditions).

In terms of wages, we find that individuals are more likely to apply for a job offering a wage close to their expectations. Going one step further, and given that in our empirical approach we control for an exhaustive array of observable characteristics, one could argue that wages in our setup are proxies for inherent *types* of both workers and firms. Thus, our results are suggestive of positive assortative matching patterns at the application stage, but more importantly, summarizes well the main distinction between search strategies of the unemployed versus the employed: while the unemployed target wage offers (types) that are very close to their own stated salary expectations, employed seekers target wages (types) which are on average above their expectations. Our reading from this is that the unemployed are trying to maximize the chances of obtaining a job offer, while the employed (performing on-the-job search) are more likely trying to climb the job ladder.

Our results also show that application decisions decline with either unemployment duration (for unemployed seekers) and job tenure (for those performing on-the-job search). This evidence is particularly useful to understand the dynamic evolution of unemployed workers over an unemployment spell, an important input for the design of unemployment insurance policies, an aspect also considered by Faberman and Kudlyak (forthcoming); the effect of tenure on job search is also relevant to understand factors behind job-to-job transitions, arguably an important mechanism to explain wage dispersion, as stated in Hornstein, Krusell, and Violante (2011)).

Our paper is related to a growing literature which use data from online job-posting/search websites in order to study different aspects of frictional markets. Kudlyak, Lkhagvasuren, and Sysuyev (2013) study how job seekers direct their applications over the span of a job search. They find some evidence on positive sorting of job seekers to job postings based on education and how this sorting worsens the longer the job seeker spends looking for a job (the individual starts applying for worse matches). Marinescu and Rathelot (2015) use information from `www.careerbuilder.com` and find that job seekers are less likely to apply to jobs that are farther away geographically. Marinescu and Wolthoff (2015) use the same job posting website to study the relationship between job titles and wages posted on job advertisements. They show that job titles explain nearly 90% of the variance of explicit wages. Gee (2015), using a large field experiment on the job posting website `www.linkedin.com`, shows that being made aware of the number of applicants for a job, increases ones own likelihood of making application.

On product markets, Lewis (2011) shows that internet seekers for used cars significantly react to posted information regarding automobiles' quality. Jolivet and Turon (2014) and Jolivet, Jullien, and Postel-Vinay (2016) use information from a major French e-commerce platform, `www.PriceMinister.com`, to study the effects of search costs and reputational issues (respectively) in product markets.

The data

We use data from `www.trabajando.com` (henceforth the website) a job search engine with presence in mostly Spanish speaking countries: as of September of 2017, the list comprises Argentina, Brazil, Colombia, Chile, Mexico, Peru, Portugal, Puerto Rico, Spain, Uruguay and Venezuela. Our data covers a sample of job postings and job seekers in the Chilean labor market, between January 1st 2008 and December 24th, 2016. The raw information in the dataset contains more than 14 million single applications, from around 1.5 million job seekers, to around 270 thousand job ads.

Our dataset has detailed information on both applicants and recruiters. First, we observe entire histories of applications from job seekers and dates of ad postings (and repostings) for recruiters. Second, we have detailed information for both sides of the market. For job seekers we observe date of birth, gender, nationality, place of residency ("comuna" and "región", akin to county and US state, respectively), marital status, years of experience, years of education, college major and name of the granting institution of the major.[4] We have codes for occupational area of the current/last job of individuals, information on their salary and

[4]This information is for any individual with some post high school education.

both their starting and ending dates.

In terms of the website's platform, job seekers can use the site for free, while firms are charged for posting ads. Job advertisements are posted for a minimum of 60 days, but firms can pay additional fees to extend this term.[5]

For each posting, we observe its required level of experience (in years), required college major (if applicable), indicators on required skills (specific, computing knowledge and/or "other") how many positions must be filled, an occupational code, geographic information ("región" only) and some limited information on the firm offering the job: its size (number of employees in brackets) and industry. Educational categories are *primary* (one to eight years of schooling), *high school* (completed high school diploma, 12 years), *technical tertiary education* (professional training after high school, usually 2-4 years), *college* (completed university degree, usually 5-6 years) and *post-graduate* (any schooling higher than college degree).

A novel feature of the dataset, compared to the rest of the literature, is that the website asks job seekers to record their expected salary, which they can then choose to show or hide from prospective employers. Recruiters are also asked to record the expected pay for the job posting, and given the same choice whether to make this information visible or not to the applicants. Naturally, one could question the reliability of wage information which will be ultimately hidden from the other side of the market. Banfi and Villena-Roldán (forthcoming) address the potential issue of "nonsensical" wage information in job ads by comparing the sample of explicit vs. implicit (job ads without any salary information) postings by firms, and find that observable characteristics predict fairly well implicit wages and vice versa. Moreover, even if employers choose to hide wage offers, they are used in filters of the website for applicant search. Hence, employers are likely to report accurately even if their wage offers are not shown because misreporting may generate adverse consequences. On the other hand, a major caveat of our dataset is the absence of information on activities performed outside the website: individuals seeking for jobs through other means, and more importantly, outcomes of job applications.

For the remainder of the paper, we restrict our sample to consider only individuals working under full-time contracts and those unemployed. We further restrict our sample to individuals aged 23 to 60. We discard individuals reporting desired net wages above 5 million

[5]As of January 3rd, 2018, the 60-day fee is CLP 69,900 + 19% VAT as posted in http://www1.trabajando.cl/empresas/noticia.cfm?noticiaid=3877, which is equivalent to USD 136 or EUR 113. There are quantity discounts to big clients, too.

pesos.[6] This amounts to approximately 8,347 USD per month[7], which represents much more than the 99th percentile of the wage distribution, according to the 2013 CASEN survey.[8] We also discard individuals who desire net wages below 159 thousand pesos (around 350 USD) a month (the legal minimum wage at the start of our considered sample). Consequently, we also restrict job postings to those offering monthly salaries within those bounds.

Our unit of analysis are individual *applications*. We restrict our sample to individuals who were actively looking for a job (i.e., made an application) and job postings that received at least one application. While we observe long histories of job search for a significant fraction of workers (some workers have used the website for several years), we consider only applications pertaining to their last job search "spell", which we define as the time window between the last modification/creation of their online curriculum vitae (cv) in the website and the time of their last submitted application or the one year mark, whichever happens first. Since individuals maintain information about their last job in their online profile, as well as contact information and salary expectations, we assume that any modification of this information is done primarily when individuals who are currently working or who have already used the website in the past are ready to search in the labor market again. We cannot infer any labor transition based on application behavior because employed individuals may keep searching for jobs, and unemployed individuals may search outside of the website. We further drop individuals who apply to more than the 99-th percentile of job applicants in terms of number of submitted applications.

Table 1 shows descriptive statistics for the job searchers in our sample. From the table we observe that the average age is 33.5 and that job seekers are comprised of mostly single males, with 59.71% being unemployed ($86,687$ unemployed seekers from a total of $215,169$ individuals.). Average experience hovers around eight years. Job seekers in our sample are more educated than the average in Chile, with 41.84% of them having a college degree, compared to 25% for the rest of the country in the comparable age group 30 to 44, according to the 2013 CASEN survey. There is also a big discrepancy by labor force status: unemployed seekers are significantly less educated in the website.

[6]In the Chilean labor market wages are usually expressed in a monthly rate net of taxes, and mandatory contributions to health (7% of monthly wage), to fully-funded private pension system (10%), to disability insurance (1.2%), and mandatory contribution to unemployment accounts(0.6%)

[7]Using the average nominal exchanges rate between 2013-16, https://si3.bcentral.cl/Siete/secure/cuadros/home.aspx.

[8]CASEN stands for "Caracterización Socio Económica" (Social and Economic Characterization), and aims to capture a representative picture of Chilean households. For data and information in Spanish, visit http://observatorio.ministeriodesarrollosocial.gob.cl/casen-multidimensional/casen/casen_2015.php

Table 1: Characteristics of Job Seekers

	Employed	Unemployed	Total
Demographics			
Age	33.77	33.25	33.46
Males	0.62	0.54	0.57
Married	0.34	0.28	0.30
Experience (years)	8.28	7.64	7.90
Wages (thousand CLP)	1,087	592	792
Tenure (weeks)	177.96	–	177.96
Unemployment duration (weeks)	–	60.20	60.20
Education level (%)			
Primary (1-8 years)	0.12	0.25	0.2
High School	17.94	36.89	29.25
Technical Tertiary	26.56	28.82	27.91
College	54.22	33.48	41.84
Post-graduate	1.17	0.55	0.8
Occupation (%)			
Management	23.5	17.85	20.12
Technology	31.59	21.21	25.39
Not declared	20.29	42.54	33.57
Rest	24.62	18.4	20.92
Search Activity			
weeks searching on website	5.24	4.83	4.99
Number of applications	1.49	1.53	1.52
Observations	86,687	128,482	215,169

From the table we can also observe that most job seekers claim occupations related to management (around 20%) and technology (around 25%) and that average expected wages are approximately (in thousands) CLP\$ $1,087$ and CLP\$ 592 for employed and unemployed seekers, respectively. For comparison, the 2013-16 average minimum monthly salary in Chile was around CLP \$ 226 thousand.[9]

In terms of search activity, the average search spell amounts to around five weeks. The amount of time searching for a job is higher for those employed than for the unemployed: 5.24 versus 4.83 weeks respectively. In terms of applications, both groups show very similar choices, with around 1.52 submitted applications.

Application probabilities and job seeker preferences

In this section, we analyze empirically which attributes of heterogeneous jobs attract more applications from heterogeneous job seekers. To do this, we first need to determine which is the relevant set of job ads for each individual in our sample. However, our dataset only

[9]The minimum wage has increased substantially in recent years. For information about the trajectory of the legal minimum wage in Chile, please see https://www.leychile.cl/Consulta/listado_n_sel?_grupo_aporte&sub=807&agr=2.

contains information on actual applications and no information is collected by the website on total number of searches nor *clicks* on job postings by individuals. Thus, we do not have sample variation in terms of job ads: we only observe those that individuals choose to apply to, but not those which are observed but then discarded by seekers. This problem of "consideration sets" (i.e. the set of products consumers are aware of) is addressed in the literature in marketing and industrial organization (Van Nierop, Bronnenberg, Paap, Wedel, and Franses, 2010; Abaluck and Adams, 2017), but our approach is essentially different. Nevertheless, we hypothesize that our network-revealed preference approach could be used in this literature, provided databases identify purchased goods by each consumer.

Market segmentation through network analysis.

We could consider the cross between all job seekers and all job ads that are time feasible in our sample, what we call the *exploded* dataset. However, there are major drawbacks from this approach: First, the *exploded* dataset makes comparisons between job seekers and job positions which may be objectively too different to consider. A typical job seeker may find more than 20,000 available job ads for her to screen and choose, implying an unrealistic effort for workers. Second, since we truly try to characterize an actual decision-making process, by introducing job ads never considered by the applicant in her choice set, there is a sample selection as we include extraneous observations into the sample. Third, a more practical issue is that the size of the estimating sample becomes simply too large to handle,[10] making the task of even simple calculations infeasible.

We could create choice sets by using clustering of job ads using their traits as in Banfi and Villena-Roldán (forthcoming). However, such an approach links workers to a fixed set of ads, with little or null cross-sectional variation across similar applicants. Instead, our approach uses revealed preferences of workers to construct individual consideration or choice sets based on coincidental choices made by other applicants. In reality, workers potentially apply to jobs considering a large number of potential characteristics, many of which we can observe. Instead of relying on our own researchers' priors about the relevance of certain characteristics to define consideration sets as, our premise is that if two individuals have a common set of applications they must be similar. We are taking whatever heuristic that applicants are using to make these decisions.

To formalize this notion, we use the network formed by job seekers to determine which

[10]With our sample constraints we have 215,000 workers who could potentially apply to 20,000 job ads when they change their CV, a very conservative lower bound if they stay actively applying for several weeks. Thus, the exploded set contains $215,000 \times 20,000 = 4,3$ billions of potential applications.

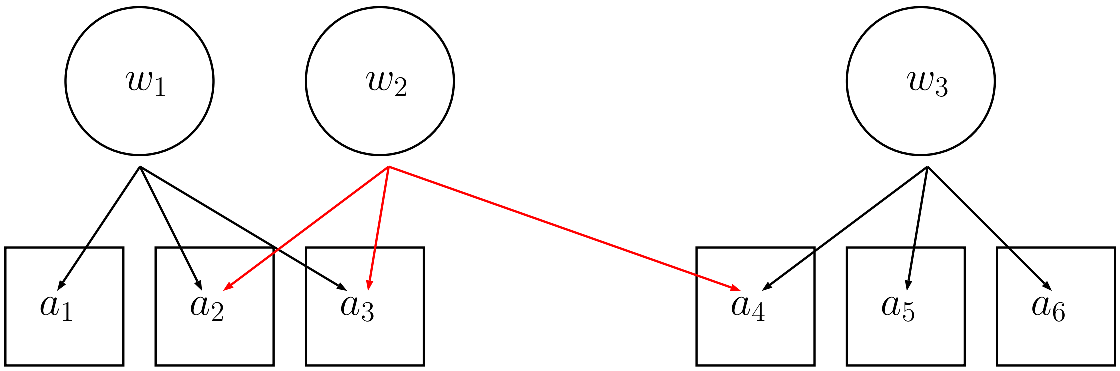

Figure 1: Example of a network formed by workers $\{w_1, w_2, w_3\}$. Worker w_1 is linked to worker w_2 by common applications to ads a_2 and a_3 but is not linked with w_3 in the network of degree 1. All workers are linked in the network of degree 2.

job postings are relevant to them. Assume that each individual represents a node in the network, and that a link between nodes is defined as *having applied to the same job posting*. For each job seeker w, we define the set of relevant job postings $\mathcal{A}_w^1$ as the union of all job postings applied by the set of all job seekers linked to w. This is what we define as a network of degree 1, since for each individual, we only consider their immediate links (1 degree of separation).

Following this logic, the network of degree 0 is the original dataset for individual w ($\mathcal{A}_w^0$), since the network contains only information of job seekers and their applications (no information on links is used). On the other hand, a network of degree 2 is defined as the network which considers both job seekers linked directly to w, in addition to those who are linked with the links of w (job seekers have 2 degrees of separation), giving rise to the set $\mathcal{A}_w^2$. We can continue with this logic iteratively, until forming the set $\mathcal{A}_w^\infty$, which is the cross between each job seeker w and all job postings a as long as they are connected somehow through the network.[11]

Figure 1 shows an example of the network algorithm and the resulting datasets. In the figure there are three workers, $\{w_1, w_2, w_3\}$ and six job postings, $\{a_1, a_2, a_3, a_4, a_5, a_6\}$. Consider worker w_1. She has applied to three jobs, thus $\mathcal{A}_{w_1}^0 = \{a_1, a_2, a_3\}$ and is linked to w_2 through applications to $\{a_2, a_3\}$. Since w_2 also applied to job position a_4, one can infer that some characteristic of a_4 is not desirable to w_1. If we consider networks of degree 1, a_4 would be included in the set of relevant ads for the first worker. Notice also that in this example, w_1 is not directly linked with w_3, or in our language, the degree of separation

[11] Technically, the set $\mathcal{A}_w^\infty$ and the exploded dataset differ if there are isolated pairs or groups of individuals who are not connected to the rest of the applicants through any ad.

between these two workers is higher than 1.

Again, considering the first worker, we have $\mathcal{A}^0_{w_1} = \{w_1, w_2, w_3\}$, and as discussed above, $\mathcal{A}^1_{w_1} = \{a_1, a_2, a_3, a_4\}$. Given that w_1 and w_2 are linked and that w_2 is linked with w_3, the relevant job ads for w_1, given a network of degree 2, is $\mathcal{A}^2_{w_1} = \{a_1, a_2, a_3, a_4, a_5, a_6\}$. In our simple example, the network of degree 2 is already the "exploded" network (all ads to all workers).

The formal definition of a one-degree-of-separation ad set for a worker w is

$$\mathcal{A}^1_w = \bigcup_{v:\mathcal{A}^0_w \cap \mathcal{A}^0_v \neq \emptyset} (\mathcal{A}^0_w \cup \mathcal{A}^0_v)$$

which can be generalized for other degrees of separation.[12] In what follows, we will concentrate on networks of degree one only.

Table 2: Number of relevant ads (a) per worker (w)

	Potential ads for a worker			Potential workers for an ad		
	All	U	E	All	U	E
percentile 10	2	2	2	1	2	3
percentile 50	16	16	19	21	25	33
percentile 90	96	104	87	233	325	385
mean	38.5	40.7	36.8	96.5	123.8	147.8
standard deviation	68.1	73.8	57.1	278.7	342.0	410.3
mean applications (%)	22.3	23.2	20.9	–	–	–

Notes: The table shows the number of relevant job postings per job seeker given a network of degree 1 (see main text). Statistics separated by labor force status of job seeker (U = unemployed, E = employed).

In table 2, we present information on the resulting number of relevant job postings per worker and workers per job posting, given a network of degrees one. The median number of relevant job postings (a) is 16 postings per job seeker, with employed seekers being related to more posts (19) than those unemployed (16). The number of potential ads exhibits quite the amount of variation, going from 2 (tenth percentile of distribution) to 104 and 89 for unemployed and employed respectively (ninetieth percentile). On the other hand, the

[12]The generalization follows a recursive definition

$$\mathcal{A}^s_w = \bigcup_{v:\mathcal{A}^{s-1}_w \cap \mathcal{A}^0_v \neq \emptyset} (\mathcal{A}^{s-1}_w \cup \mathcal{A}^0_v)$$

which depends on $\mathcal{A}^0_w$ and the definition of $\mathcal{A}^1_w$.

median number of workers related to each job advert is 233, with employed individuals being attached to more job ads than those unemployed. Given the sets of related job ads, mean application rates,[13] are 22.3% for the entire sample, with unemployed seekers applying to 23.2%, while employed ones do so for 20.9% of their relevant ads.

Although our network approach allows us to build choice sets for applicants, it is unlikely that all the ads in the set are equally considered. In any given network induced set of choices for each worker w, there is heterogeneity in the relevance of job ads, according to how strong the link between two workers is. Intuitively, the bigger the overlap in submission choices by both workers, the closer they are and the more relevant the additional job ads are for each other. As an example, consider worker w_2 in figure 1. Since w_2 and w_1 submit common applications to several common positions, they must have similar preferences and qualifications. Then, the likelihood that w_2 truly considers applying to job ads to which w_1 applied to must be high. In contrast, w_2 and worker w_3 share less applications, so the likelihood that w_2 considered $\{w_5, w_6\}$ is lower.

To give more formality to this intuition, we construct a weight function $q(w, a)$ for each worker w and job position a. We start with function $b(w, v)$ which we apply to all pairs of linked workers w and v, as a measure of how similar they are in terms of application decisions. We construct b, given some general restrictions:

1. $b(w, v) \in [0, 1]$
2. $b(w, w) = 1$
3. $b(w, v) = 0$ if and only if $\mathcal{A}^0_w \cap \mathcal{A}^0_v = \emptyset$

On top of conditions 1-3 above, we want the function $b(w, v)$ to be monotonic in set similarity. A particular functional form that satisfies these conditions is

$$b(w, v) \equiv \frac{|\mathcal{A}^0_w \cap \mathcal{A}^0_v|}{|\mathcal{A}^0_w \cup \mathcal{A}^0_v|} \tag{1}$$

where $|S|$ is the cardinality (number of elements) of set S. Equation (1) is also known as the *Jaccard Similarity Index* between two groups (Jaccard, 1901).

[13] Defined as the number of effective applications to total ads for worker w:

$$\frac{|\mathcal{A}^0_w|}{|\mathcal{A}^1_w|}$$

We then define the weight of an ad a for a worker w as:

$$q(w,a) = \max_{v:a\in\mathcal{A}_v^0} \{b(w,v)\} \tag{2}$$

Intuitively, we consider the importance of a particular job ad a in the choice set of w given the similarity between the choice set of w and the most similar choice set of any other applicant v, linked to w. It is easily verified that given this proposed weighting function:

1. $q(w,a) = 1$ if and only if $a \in \mathcal{A}_w^0$. (w applies to a)
2. $q(w,a) \in [0,1)$ if and only if $a \notin \mathcal{A}_w^0$ (w does not apply to a)
3. $q(w,a) = 0$ if and only if $a \notin \mathcal{A}_w^1$ (a is not in the choice set of w)

These definitions operationalize a similarity notion between workers and ads we use to estimate application equations.

Estimation of the application equation

For the constructed dataset, we estimate preferences of job seekers, based on their observed characteristics along the ones posted by ads which are relevant to them. More specifically, we estimate a linear regression of the form

$$y_{aw} = X_{aw}\beta_{aw} + \sum_{k_c}\sum_{p=1}^{\bar{P}}\{\beta_{k_o p}(z_{k_o})^p\} + \beta_{k_o}z_{k_o} + \sum_{k}\sum_{\ell}\mathbf{1}_{\{k\neq\ell\}}\beta_{k\ell}z_k z_\ell + \epsilon_{aw} \tag{3}$$

where y_{aw} is a dummy variable that takes the value of one if a job seeker w applies to posting a, and zero otherwise. In X_{aw}, we control for observed job and worker characteristics, which do not overlap. The list of variables for the job includes firm size, dummies for firm industry (1 digit) and specific job requirements (computer knowledge, or some other form of specific knowledge) and controls for specific job characteristics: type of contract (full/part time), number of vacancies needed to be filled and controls for job titles.[14] For individuals, we control for marital status (dummy variable for marriage), gender (dummy for male), an interaction between married and males and quintic polynomials for the age of the job seeker and for the amount of time (measured in weeks) in either the current job (for those employed) or in unemployment (for unemployed seekers). For both seekers and ads, we include a variable of whether the wage expectation (for seekers) or the wage expected to be paid (for jobs) is made explicit or not. As control for business cycle conditions, we

[14] We follow work by Marinescu and Wolthoff (2015) and Banfi and Villena-Roldán (forthcoming).

include the national unemployment rate of the Chilean economy during the date (quarter) in which the application took place.[15] The effects of these characteristics impact the level of the probability of application, and therefore are related to the *intensive margin* of the application process that have been more profusely studied in the literature (Faberman and Kudlyak, forthcoming; Gomme and Lkhagvasuren, 2015; Leyva, 2018; Mukoyama, Patterson, and Şahin, 2018)

On the other hand, to provide a novel measurement of the *selective margin* of job search, we include a set of controls for the *misalignment* (which we denote by z) between characteristics required by firms vs. the characteristics of the job seeker. For continuous variables, which we denote by k_c, and encompass the level of education, years of experience, log wages and regional distance[16], we define z_{k_c} as the simple difference between the value of the characteristic required by the position and value of the characteristic possessed by the job seeker.

For occupations, the variable z_{k_o} is defined as a dummy that takes the value of one when the category in the job posting is different from the characteristic of the worker and zero when they are the same.

In equation (3), for each of the continuous dimensions k_c we include in the regression a polynomial of order $P = 5$ to assess whether non-linearities exist in the effect of these *misalignments* on application decisions. In this way, we capture if *over*-qualified ($z_{k_c} < 0$) jobseekers behave differently from *under*-qualified ($z_{k_c} > 0$) ones. We estimate the above equation separating our sample between the employed and unemployed to assess whether on-the-job search differs from unemployed search behavior. Finally, we also consider interaction effects between different *misalignment* levels and weight worker-ad observations by q described above.

Results on the intensive margin: Applicants and ad traits

Table 3 shows coefficients multiplied by 100 from the estimating equation (3) using ordinary least squares. We report estimates by employment status and whether we perform network weighting or not of our estimates. Results related to polynomials on continuous *misalignment* variables are presented later.

In terms of demographics and family composition, male workers apply more than females.

[15] For worker-ad pairs that are matched given our network algorithm, the date of an actual application does NOT exist. In those cases, we impute the date of application by the mode date of applications of the linked workers to that particular job ad.

[16] We only observe the region where employers and applicants are located, and we compute distance as hundreds of kilometers between regional capital cities

Table 3: Intensive margin coefficients by labor status

VARIABLES	(1) Employed weights	(2) Employed no weights	(3) Unemployed weights	(4) Unemployed no weights
Married	3.979***	-0.373	2.495***	-0.302*
	(0.493)	(0.327)	(0.316)	(0.181)
Male	1.316***	0.252***	1.493***	0.278***
	(0.063)	(0.032)	(0.053)	(0.025)
Male x Married	-1.228***	-0.037	1.107***	0.436***
	(0.106)	(0.057)	(0.098)	(0.049)
Explicit wage (w)	-0.019	-0.113***	-0.035	-0.041**
	(0.050)	(0.025)	(0.043)	(0.021)
Explicit wage (a)	-1.862***	-0.644***	0.521***	0.064**
	(0.080)	(0.039)	(0.059)	(0.028)
No. of Vacancies (a)	-0.008***	-0.003	0.030***	0.010***
	(0.001)	(0.002)	(0.001)	(0.001)
Ad duration (weeks)	0.024***	0.001	-0.184***	-0.026***
	(0.003)	(0.001)	(0.003)	(0.001)
Observations	2,955,376	2,955,376	4,330,259	4,330,259
R-squared	0.144	0.040	0.140	0.040

Notes: Regression coefficients from a linear regression on application decisions. Dependent variable is y_{aw}, a dummy for the existence of a job application. Each regression controls also for polynomials and interactions in *misalignment* as well as age of the worker, firm size, contract type, dummies for different types of requirements of the job and characteristics of the firm (see details in the main text). Standard errors in parentheses. One, two, and three asterisks indicate significance at 10%,5%, and 1%, respectively.

We find that employed married individuals apply less than non-married counterparts. The effect of the interaction term of being married and a male job seeker varies on the labor status: the coefficient is positive and significant for the unemployed, but negative and insignificant for the employed. Hence, unemployed married males tend to apply more intensively than other groups, suggesting a higher opportunity cost. Although Faberman and Kudlyak (forthcoming) find that males apply less than females, the results are not directly comparable as we condition on a rich set of job titles and other observables that are unavailable in their data. Although our search effort measures differ, Aguiar, Hurst, and Karabarbounis (2013), in line with us, find that males and singles spend more time searching for jobs in comparison to females and married counterparts. Van Hooft, Born, Taris, and van der Flier (2005) attribute lower intensity of job search by married due to a higher cost to family reallocation.

Table 4: Intensive margin coefficients by gender and labor status

VARIABLES	(1) Employed female	(2) Employed male	(3) Unemployed female	(4) Unemployed male
Married	1.131*	10.240***	3.759***	-2.754***
	(0.681)	(0.808)	(0.403)	(0.589)
Explicit wage (w)	-0.252***	0.230***	-0.432***	0.438***
	(0.085)	(0.062)	(0.065)	(0.058)
Explicit wage (a)	-2.112***	-1.578***	1.216***	-0.280***
	(0.131)	(0.102)	(0.088)	(0.079)
No. of Vacancies (a)	0.009***	-0.039***	0.044***	0.003**
	(0.002)	(0.001)	(0.001)	(0.001)
Ad duration (weeks)	0.008	0.024***	-0.168***	-0.205***
	(0.005)	(0.004)	(0.004)	(0.004)
Observations	1,054,600	1,900,776	1,929,679	2,400,580
R-squared	0.149	0.154	0.146	0.153

Notes: Regression coefficients from a linear regression on application decisions. Dependent variable is y_{aw}, a dummy for the existence of a job application. Each regression controls also for polynomials and interactions in *misalignment* as well as age of the worker, firm size, contract type, dummies for different types of requirements of the job and characteristics of the firm (see details in the main text). Standard errors in parentheses. One, two, and three asterisks indicate significance at 10%,5%, and 1%, respectively.

Our results also show that individuals who choose to be explicit about their wage expectations at the time of an application are more likely to apply to any position, which is true for both unemployed and employed seekers, although the effect is somewhat stronger for the employed sample. While there are no theories to explain wage expectation disclosure in the literature, we conjecture that workers who have conveyed this information to employers apply more to compensate potentially lower call back rates.

The table also shows that an explicit wage in the job ad affects negatively the decision to apply for for all type of job seekers, but the effect is less important for the unemployed. This is consistent with findings in Banfi and Villena-Roldán (forthcoming) who show that ads with hidden wages tend to attract more applicants due to a higher likelihood of potential wage flexibility or bargaining, as suggest in the Michelacci and Suarez (2006) model. Moreover, unemployed individuals apply slightly more to job ads that advertise higher number of vacancies to be posted (just 0.05% higher chance to apply for an extra vacancy), while the employed do not significatively react to that information. The small response to a marginally higher likelihood of receiving an offer suggests an important role for recruiting selection on the employer side, i.e. non-sequential employer search van Ours and Ridder (1992); van Ommeren and Russo (2013).

The effect of the perceived "age" of the job ad has also different effects depending on the labor force status of the individual: unemployed seekers seem to have a distaste for job ads that are older (in weeks), while those employed prefer them at the margin. The negative effect of for the unemployed seems related to stock-flow matching behavior[17]: new job seekers in the website (the flow) apply to the stock of job ads. When time passes, the inflow of job seekers becomes part of the stock of individuals, who then try to match with the new flow of job positions, as suggested by evidence in Gregg and Petrongolo (2005) and Coles and Petrongolo (2008). Our results are also consistent with applicants reacting to "phantom" ads, which may be filled positions by the time of the potential application, as in Albrecht, Decreuse, and Vroman (2017) and Chéron and Decreuse (2016). Nevertheless, we know of no model implying different applicants' reactions to elapsed posting duration according to their employment status.

Results when we do not weight applications by the function $q(w, a)$ are shown in columns 2 (for employed) and 4 (for unemployed) of table **??**. Almost all estimates have the same sign as in the weighted regressions, but the magnitudes seem attenuated. A possible interpretation is that our weighing method reduce the importance of irrelevant job ads into the sample, reducing the inevitable inclusion of extraneous job ads into applicants'consideration sets.

In table 4, we run weighted the regressions by gender. A striking finding is that married females apply substantially less to ads if employed, which may have strong implications for climbing the job ladder, an important source of wage growth. Being married generally reduces application probabilities, except for unemployed males. There are no important gender differences in application behavior once wage expectations have been disclosed. However, females are more reluctant to apply to jobs with explicit wages, perhaps due to a higher valuation for flexible job conditions, as shown in Wiswall and Zafar (2018).

Gender shapes the application response to an extra vacancy in job ads: unemployed females have substantially higher positive response, while employed males reduce their application likelihood. Employed males response to ad "age" is much higher than employed females. On the other hand, the distaste for old postings is similar for both genders.

[17] References are Taylor (1995); Coles and Muthoo (1998); Coles and Smith (1998); Ebrahimy and Shimer (2010)

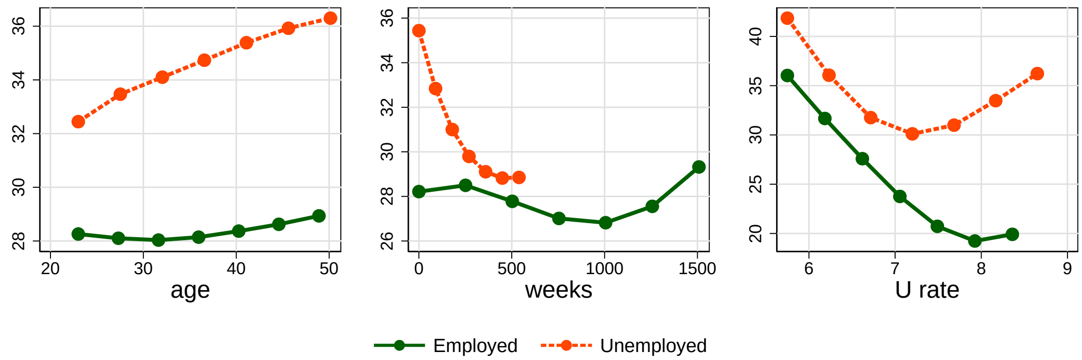

Figure 2: Predicted application probabilities for different ages, number of weeks in the current labor force status, and national unemployment rate at the time of the application decision, given results from equation (3). The figure is computed using the coefficients associated to a polynomial of order 5 on each variable and leaving the rest of regressors at their sample mean.

Results on the intensive margin: Life-cycle, duration and business cycle effects

We report the predicted application probability varying age, duration of employment status, and unemployment while keeping the other covariates at their mean values in figure 2.

In the left panel, we first observe that unemployed apply more at all ages. For the employed, applications decrease with age, while for the unemployed, there is a hump-shaped profile. For both types of seekers, job applications decrease from early on (around age 27). This evidence is consistent with findings in Choi, Janiak, and Villena-Roldán (2015) and Menzio, Telyukova, and Visschers (2016), among others, with respect to job finding rates and employment to employment transitions over the life-cycle in the US. For the Chilean labor market, Naudon and Pérez (2018) obtain similar results.

The middle panel in the figure shows a decreasing application probability as the search duration increases, measured as the time elapsed between the finishing date of the previous job and the application date. The extended range of durations (from percentiles 5 to 95 of the distribution) suggests that equalizing traditional unemployment duration with our measure of search duration is far-fetched. Thus, an appropriate interpretation is that individuals who have lost jobs and are website users concentrate their applications soon after the separation. In contrast, workers with more tenure increase their likelihood of being on-the-job searchers in the website.

In terms of business cycle conditions, the right panel of figure 2 shows a nonlinear relationship between the unemployment rate, our cyclical variable, and application decisions. As expected, the unemployed apply more than their employed counterparts. For both labor status, the application probability reaches a minimum around a 6.6% unemployment rate. An increment from 7% to 8% of unemployment raises the application likelihood by 15% for both employed and unemployed groups, a very large effect. As the unemployment rate approaches to 9%, the unemployed apply to more than half of the relevant ads, while the unemployed are slightly below. For a unemployment rate below 6.6%, all workers also increase the likelihood of applying. The increasing part of the curve suggests that search effort tries to offset the scarcity of available jobs when the unemployment is high, i.e. search effort is countercyclical, in line with Faberman and Kudlyak (forthcoming) and Mukoyama, Patterson, and Şahin (2018). However, the decreasing part of the curve shows a procyclical pattern of search effort, as advocated by (Gomme and Lkhagvasuren, 2015). Yet (Leyva, 2018) finds roughly acyclical search effort. Our finding of non-monotonicity of the effect helps reconciling these heterogeneous pieces of evidence in the literature. Besides this point, it is quite surprising that the cyclical response of applications is large and similar for employed and unemployed seekers. This suggests that on-the-job searchers react to the same incentives that spur the unemployed: search harder when facing higher unemployment risk in downturns, or more abundant job opportunities in tight labor markets.

Selective Margin: Misalignment and applications.

We present the effect of *misalignment* in continuous dimensions (education, experience, log wages, and distance). As noted above, for each dimension we take the simple difference between what is required in the job ad (years of experience, for example) and what the job seeker possesses. Then, a negative value for this *misalignment*measure means that the individual is "overqualified" in the particular dimension and that the individual is "underqualified" if it is positive. For instance, if an applicant has more years of experience than the minimum required in the job ad, she is "overqualified" in terms of experience. With some abuse of language, we will define that a worker applying to an ad posting a wage lower than her own expectation is "overqualified".

In figure 3 we present graphically results of the effect of *misalignment*in years of education, years of experience, log wages, and regional distance on application decisions. The figure shows predicted application probabilities ($\hat{y}_{aw}$ from the estimates of equation 3), when a particular continuous dimension mistmatch*misalignment*(z_{k_c}) varies, keeping all other observables at their sample mean, including the *misalignment* in other dimensions. Given that each *misalignment*dimension enters the equation as a fifth-order polynomial and that there are interactions between them, the computed effect is highly non-linear and depends on which value the other control variables take. The considered range for z_{k_c} is bounded by its

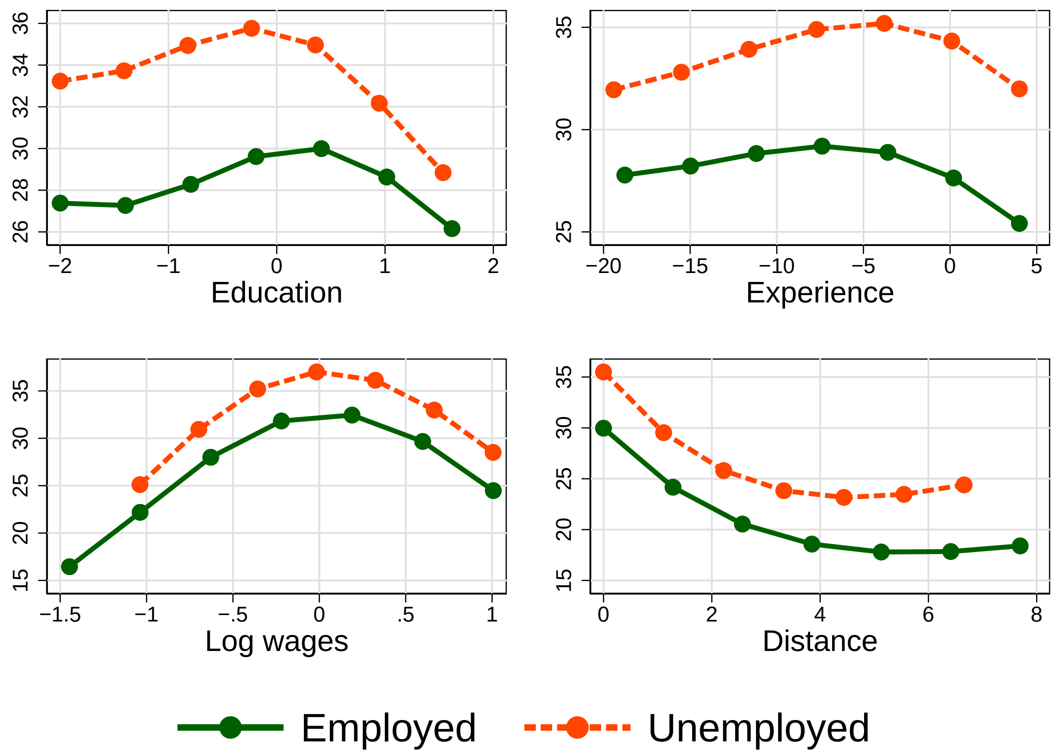

Figure 3: Predicted application probabilities, given results from eq. (3) and different levels of misalignment in the selected variable x (see main text for details). The rest of regressors are at their sample means.

5th and 95th percentiles.

As seen in figure 3, predicted application probabilities are always higher for the unemployed. Moreover, job seekers in both labor market states tend to align themselves with the advertised requirements of job postings. This is represented by an inverted U-shaped relationship between *misalignment* and application probability (all else constant) for education, experience, log wages, and by a decreasing line in the case regional distance (except for jobs located more than 500 kilometers away).

The application probability for the employed roughly peaks at 0.5 years of underqualification; that is, on-the-job applicants are more likely to aim for job ads asking more education than they have. In contrast, the unemployed curve for education peaks below zero, implying that unemployed workers have a slight tendency for applying to job ads for which they are overqualified in education. Moreover, there are asymmetric responses: there is a steeper decline in application probabilities to the left of the peak than to the right of it for the unemployed. These findings suggest that employed seekers seem more ambitious, assuming that jobs requiring more education are better. Thus , on-the-job search has a job ladder motive.

The experience dimension curves peak below -5 and becomes roughly flat to the left. This means that job seekers tend to have more than five years of experience than the minimum required, and do not refrain from applying much if they are even more overqualified in experience. The main reason for the average *misalignment* in this dimension is that our sample are attached to the labor force, with significant number of years of experience. The application probability curves for both labor status are roughly parallel, with just a mild steeper decline for the unemployed.

The plot at the lower-left panel reveals that differences in log wages greatly affect application probabilities: for unemployed seekers, the relative probabilities fluctuate between 15% and 30%, while for employed seekers, the range is wider, from around 10% to 23%. Given that our estimates control for all other observables across job positions and job seekers, and that the regression controls for interactions, we can interpret the *misalignment* in log-wages as *misalignment* in job and worker unobserved productivities. Controlling for all observables, higher paying jobs and job seekers with higher earnings expectations must be of higher skill on average, and viceversa.

The curve for the employed lies below the one of the unemployed and peaks at a higher misalignment level, slightly above zero. These facts portray on-the-job searchers as more adventurous than their unemployed counterparts, which could be spurred by the better outside options of the former. Thus, unemployed seekers seem more conservative and try to maximize the probability of getting hired, while on-the-job searchers are interested in climbing the job ladder.

The lower-right panel depicts the predicted probability as a function of the distance between the regional capital of the applicant and the regional capital of the job in hundreds of kilometers. For ads located relatively close to the applicants, the likelihood of application decreases quite fast. The probability curves for the employed reaches a minimum for jobs located 400 kilometers away from the applicant, while for higher distances a tenuous increase is apparent. The unemployed curve reaches its minimum for a distance slightly less to the left. Marinescu and Rathelot (2018) and Manning and Petrongolo (2017) estimates imply a much larger drop in the likelihood for applying to jobs as distance increments, although our estimates are not directly comparable in that we control for a substantially richer set of variables.

Selective Margin and Time Variation.

A large literature on studying how job search varies over time in different domains. Life-cycle is stressed for early human capital accumulation, fertility and retirement decisions. Unemployment and employment durations matter for unemployment insurance and job-ladder climbing decisions. The business cycle domain is quite important to understand the overall impact of job search effort on the unemployment rate, job allocation and productivity. Thus, we naturally ask how the effect *misalignment*in different dimensions at the time of searching for a job, interact with time in these domains. This could be seen as studying the interaction of the effects studied separately in the the two previous sections. Below, we separate estimation samples by quartiles of the three time variables: age of the worker, weeks in the current labor market status, and level of aggregate unemployment at the time of the application decision. After estimating the regression in each of these sub-samples, we repeat the exercise in figure 3 of producing application probabilities by levels of *misalignment*. In the figures below, Q_1 to Q_4 represent the quartiles in ascending order, while the first row (group of three panels) shows effects for the employed, while the second row, for the unemployed.

In figure 4 we report the results of the exercise for the education dimension. In terms of age effects for the employed, the upper left panel (labeled *by Q of age*) show that applications occur more frequently for the younger (Q1) and the older (Q4) groups, but the inverted U of the latter is much flatter. Q1 curve also peaks to the left of the other groups. These facts indicates that young searchers comply more to education ad requirements compared to Q4. In contrast, the middle groups Q2 and Q3 exhibit very similar profiles with overall higher

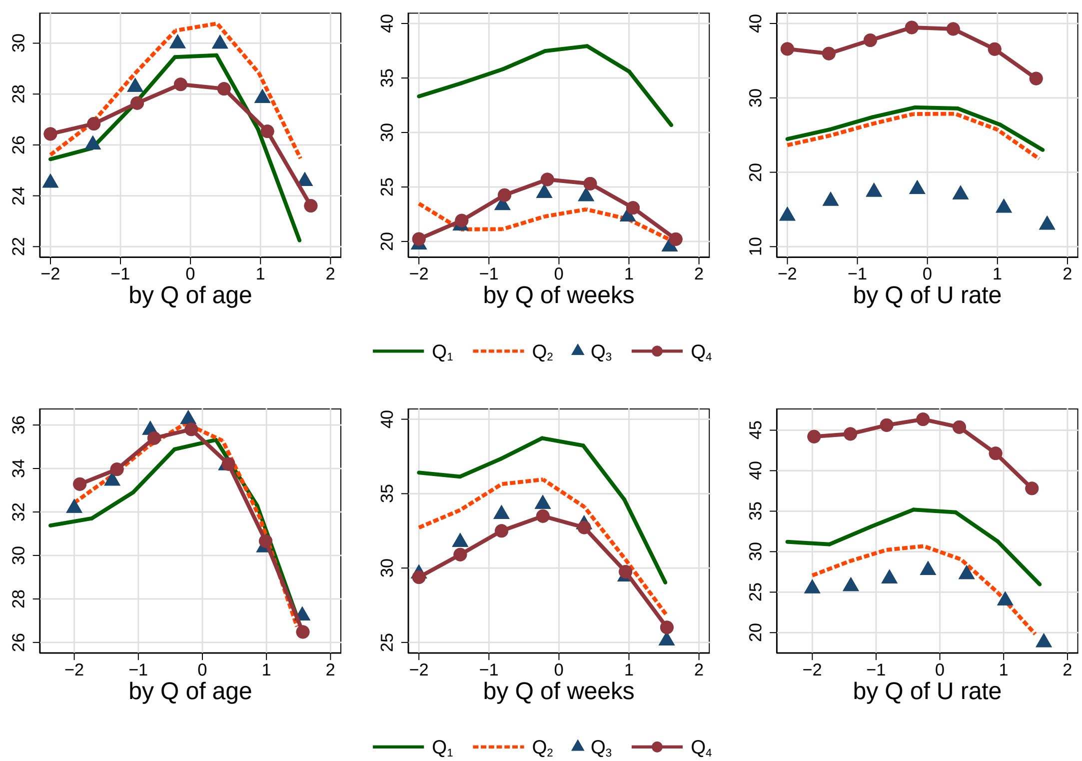

Figure 4: Predicted application probabilities for EMPLOYED (first row) and UNEMPLOYED (second row), given results from eq. (3) and different levels of misalignment in EDUCATION (see main text for details). The rest of regressors are at their sample means. The figure shows results when we split the sample in quartiles of age, weeks in the current labor force status, and the national unemployment rate.

application probability than Q1 and Q4. As for the unemployed, we observe in the lower-left panel that Q1 is different from the rest: the application probability is lower and decays a bit more sharply than that of Q2, Q3 and Q4. In contrast to the employed case, younger workers do not comply more to education ad requirements compared to others quartiles of the unemployed.

In line with findings of figure 2, the upper center panel of figure 4, the application probability increases as the tenure of the applicants in their current jobs increases, especially for Q3 and Q4. The inverted U shapes peak in a positive *misalignment*level for all but Q4, suggesting a job ladder motive of seekers applying for jobs for which they are underqualified in terms of education. For Q4, however, the curve peaks slightly to the left of zero, suggesting factors other that ladder climbing are at stake. For the unemployed, the lower center panel of figure 4 shows a roughly parallel downward displacement of the application probability curves as quartiles contain longer unemployment durations. . All curves peak slightly to the left of zero, suggesting a conservative searching strategy that probably intends to secure a job.

[COMPLETE LATER: STRANGE RESULT] The upper right panel depicts the application probability of the employed curves as the unemployment rate increases. The scale of the panel reveals that a higher unemployment remarkably affects application probabilities in a non-monotone way. The Q3 group displays the lowest likelihood, averaging around 15%, while Q1 and Q4 are on top nearly 30%.

The case of *misalignment*in experience is displayed in figure 5. In the left two panels we observe that there are significant effects of age on how individuals align themselves with experience requirements of job ads: older workers exhibit increasing flatter curves in those panels, which implies that they do not align themselves to any particular level of required experience. Since older applicants have high experience, that information has little value to make application decisions.

The effect of duration (two center panels) is different for employed (top panel) than for the unemployed (lower panel). For the earlier, our results show that longer tenure times have the effect of flattening out the application curve, so individuals performing on-the-job search become less sensitive to experience requirements the longer they have been employed in their current job. For the latter, the opposite occurs: unemployed seekers who have been longer in unemployment exhibit a slimmer application bell curve.

With respect to economy wide conditions, effects are also different depending on labor force status. For the employed (top right panel), higher unemployment rates are associated

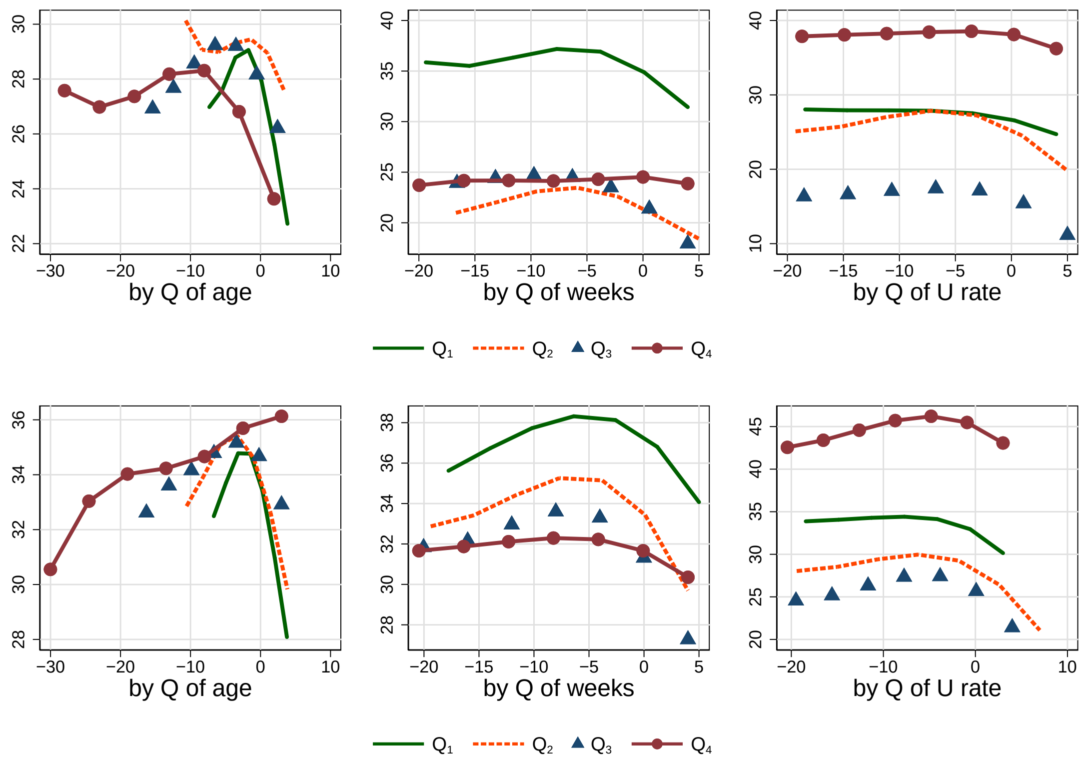

Figure 5: Predicted application probabilities for EMPLOYED (first row) and UNEMPLOYED (second row), given results from eq. (3) and different levels of misalignment in EXPERIENCE (see main text for details). The rest of regressors are at their sample means. The figure shows results when we split the sample in quartiles of age, weeks in the current labor force status and the national unemployment rate.

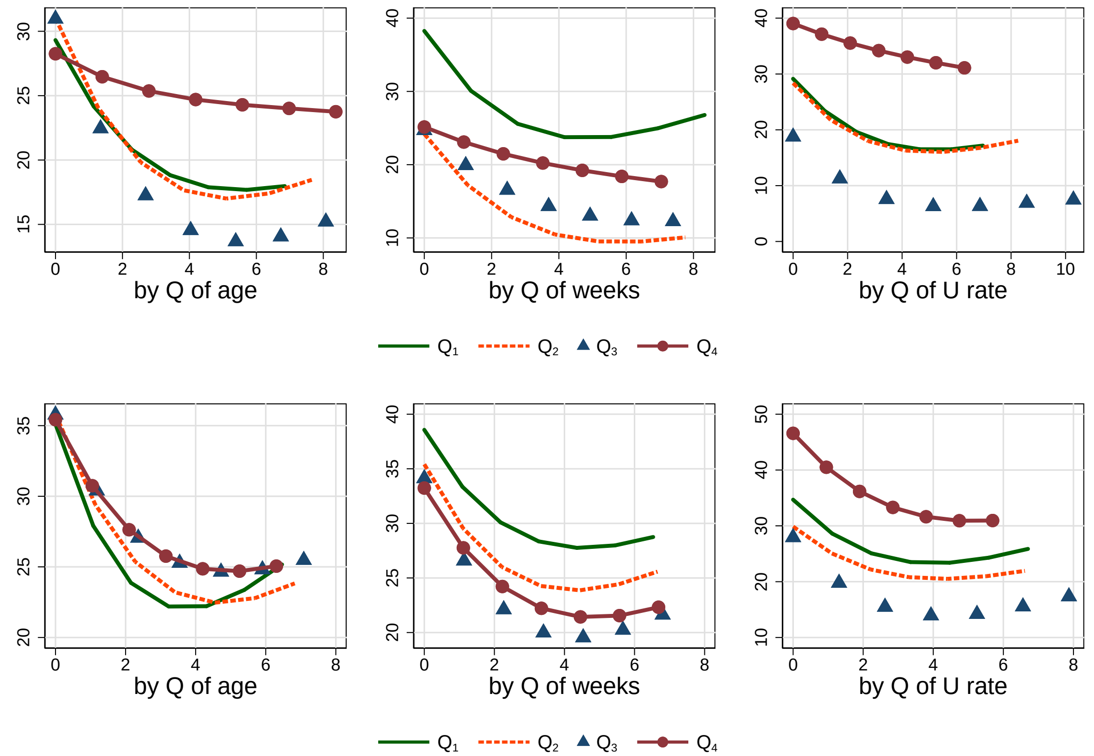

Figure 6: Predicted application probabilities for EMPLOYED (first row) and UNEMPLOYED (second row), given results from eq. (3) and different levels of misalignment in DISTANCE (see main text for details). The rest of regressors are at their sample means. The figure shows results when we split the sample in quartiles of age, weeks in the current labor force status, and the national unemployment rate.

with comparatively lower application probabilities for the extreme cases of being over qualified in the experience dimension: the left part of the application curve is lower under Q_4 than for the rest. For the unemployed (lower right panel). the effect is more marked for the extreme cases of under qualification of workers with respect to job ads.

In figure 6 we present the results for regional *misalignment*. The interaction between life-cycle and duration effects with *misalignment*in the location dimension are not salient, as reflected by the closeness of lines related to different quartiles in the left and center panels of the figure. However, there are significant effects with respect to aggregate conditions. In both top and lower right panels of the figure, we can observe that application probabilities fall significantly when the economy is experiencing high levels of unemployment, decreasing by almost a half. This is true for both the employed and the unemployed, but the effects are slightly stronger for the latter.

Finally, figure 7 shows results when we consider the (log) wage dimension. As stated

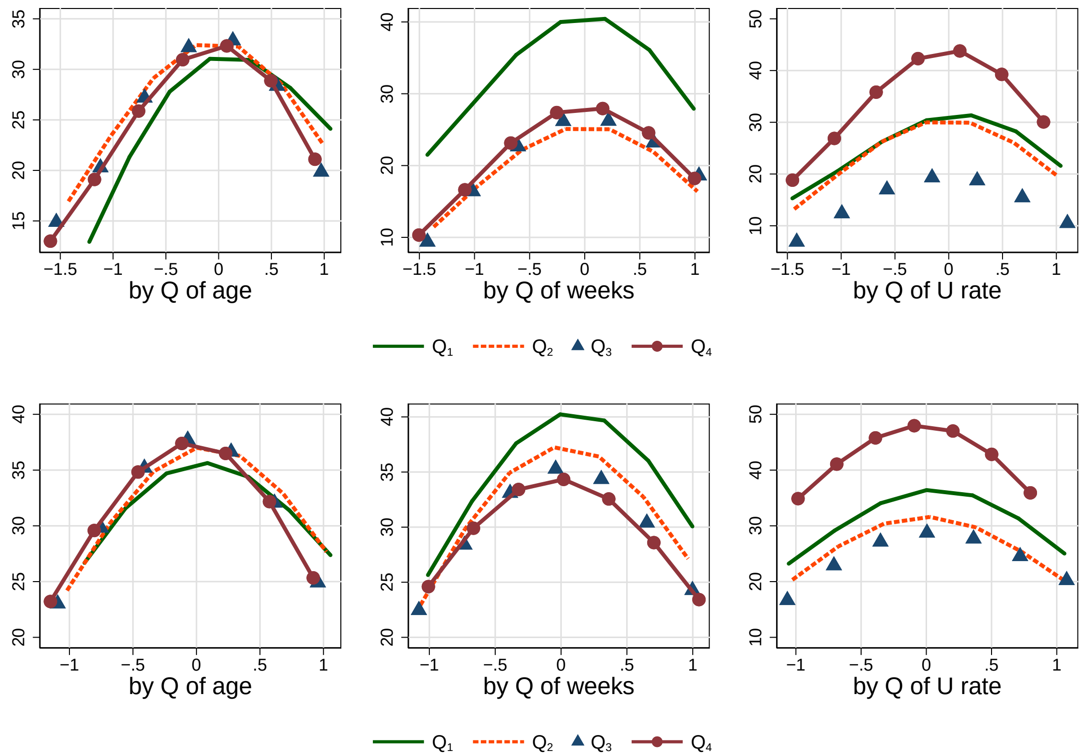

Figure 7: Predicted application probabilities for EMPLOYED (first row) and UNEMPLOYED (second row), given results from eq. (3) and different levels of misalignment in LOG-WAGES (see main text for details). The rest of regressors are at their sample means. The figure shows results when we split the sample in quartiles of age, weeks in the current labor force status, and the national unemployment rate.

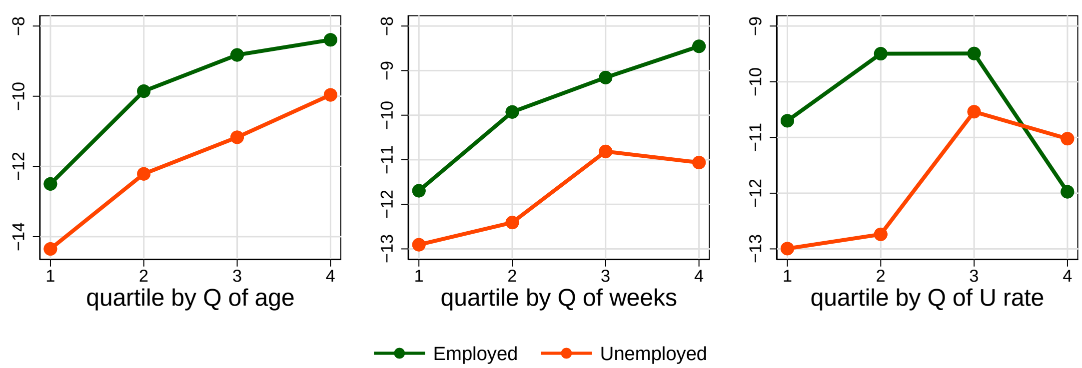

Figure 8: Predicted effect of applying to an ad with occupation mismatch for EMPLOYED and UNEMPLOYED given results from eq. (3). The other regressors are at their sample means. The figure shows results when we split the sample in quartiles of age, weeks in the current labor force status and national unemployment rate.

before, one can view *misalignment*in this dimension as evidence of overall sorting of workers and job positions along types.

Results from the figure show small effects of either life-cycle, duration or business cycle conditions on application decisions, when individuals are considering *misalignment*in the log-wage dimension. Since we do find some effects above in terms of the other dimensions interacted with age, duration and unemployment rates, the lack of shifts in the application curves in figure 7 may reflect some averaging of all the effects, which in the end imply small total effects when individuals consider *misalignment*in log-wages.

Conclusions

Using data from a Chilean job posting website, in this paper we uncover several facts regarding the nature of job search in an online setting. Given our unique setup, we can deconstruct behavior into two dimensions: an *extensive* dimension (number of applications sent) and an *intensive* one, where we analyze what affects the decision of sending an application at the margin.

For both dimensions, we find marked heterogeneity across those looking from unemployment and those performing on-the-job search, specially along the gender dimension. We find discouragement effects, in the sense that job search efforts decay with the time spent in the current labor force status (unemployment and employment). We also find new evidence regarding job search behavior during the business cycle: while the number of applications sent by individuals is clearly counter-cyclical (higher number when aggregate unemployment is down), aggregate conditions have a non-linear effect on the decisions of sending additional applications.

In the last section, we show how job seekers react to *misalignment* in key dimensions between own characteristics and characteristics required by job postings (level of education, years of experience, location, required occupation and log-wages) and find that there is significant alignment between requirements and characteristics. This alignment is affected by life-cycle, duration (time spent in the current labor force status) and aggregate business cycle effects.

References

Abaluck, J., and A. Adams (2017): "What do consumers consider before they choose? Identification from asymmetric demand responses," Discussion paper, National Bureau of Economic Research.

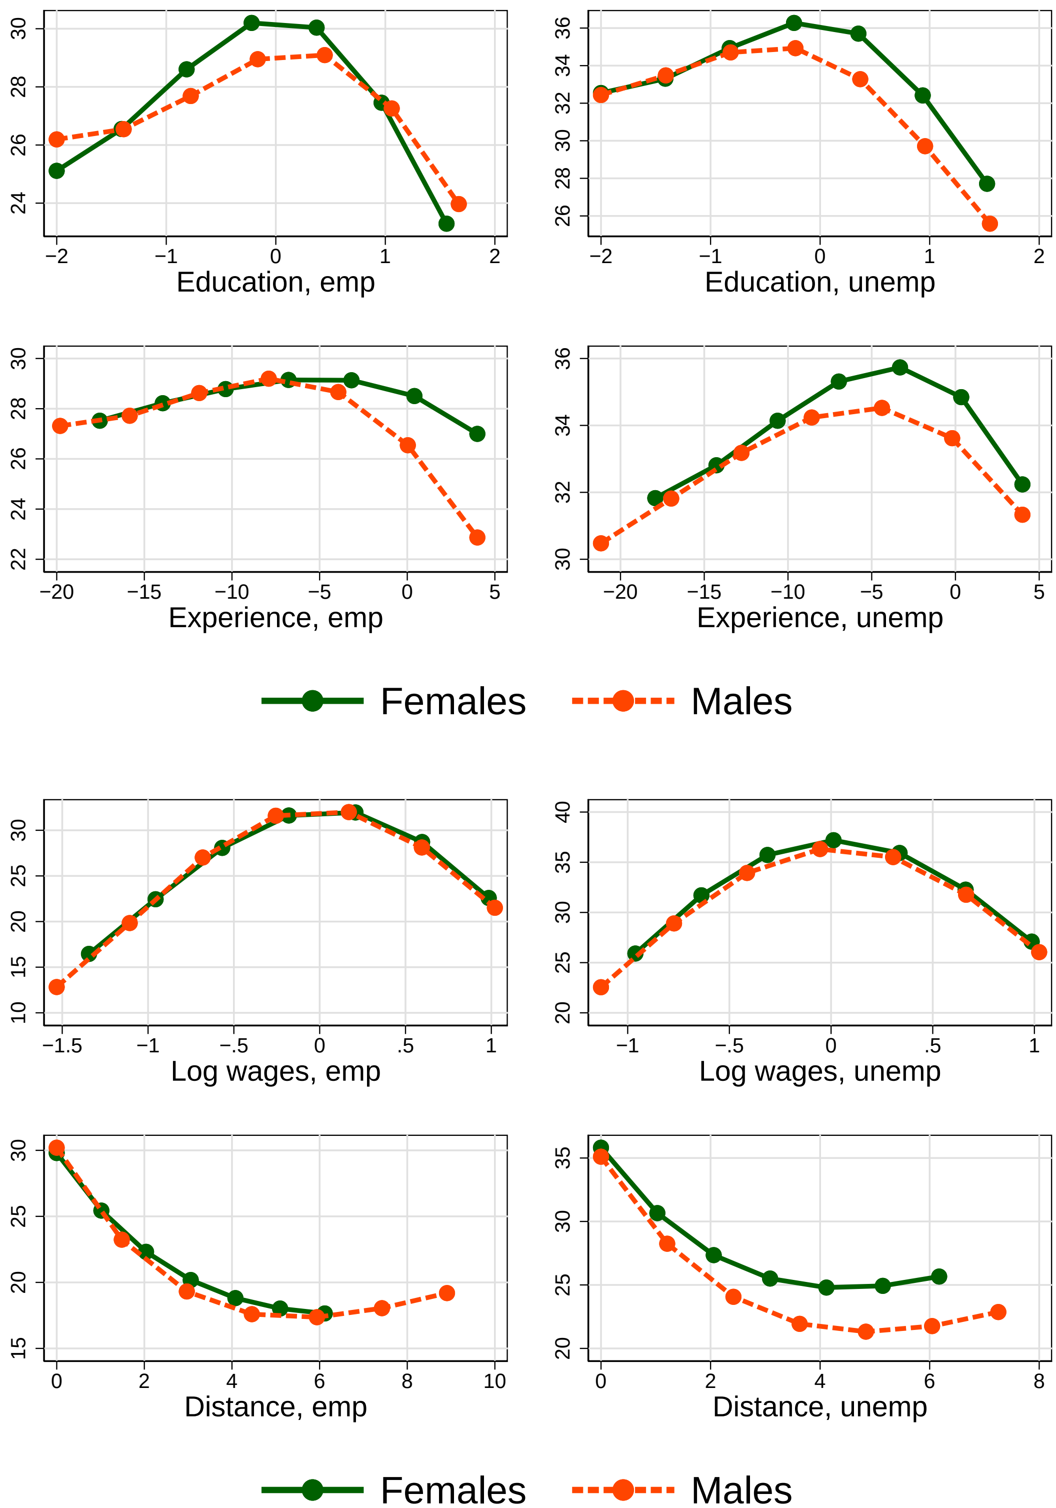

Figure 9: Predicted application probabilities, given results from equation (3) by employment status and gender at different levels of misalignment in selected variables: education, experience, log wages, and distance (see main text for details). The rest of regressors are at their sample means.

Figure 10: Predicted effect on application probabilities, given results from equation (3) by employment status and gender. The rest of regressors are at their sample means.

Aguiar, M., E. Hurst, and L. Karabarbounis (2013): "The life-cycle profile of time spent on job search," *American Economic Review*, 103(3), 111–16.

Albrecht, J., B. Decreuse, and S. Vroman (2017): "Directed search with phantom vacancies," Working paper, Aix-Marseille School of Economics.

Banfi, S., and B. Villena-Roldán (forthcoming): "Do High-Wage Jobs Attract more Applicants? Directed Search Evidence from the Online Labor Market," *Journal of Labor Economics*.

Carrillo-Tudela, C., and L. Visschers (2014): "Unemployment and Endogenous Reallocation over the Business Cycle," Cesifo working paper 4079.

Chéron, A., and B. Decreuse (2016): "Matching with Phantoms," *Review of Economic Studies*, (0), 1–30.

Choi, S., A. Janiak, and B. Villena-Roldán (2015): "Unemployment, Participation and Worker Flows Over the Life-Cycle," *The Economic Journal*, 125(589), 1705–1733.

Coles, M., and B. Petrongolo (2008): "A Test between Stock-Flow Matching and the Random Matching Function Approach," *International Economic Review*, 49(4), 1113–1141.

Coles, M. G., and A. Muthoo (1998): "Strategic Bargaining and Competitive Bidding in a Dynamic Market Equilibrium," *Review of Economic Studies*, 65(2), 235–260.

Coles, M. G., and E. Smith (1998): "Marketplaces and Matching," *International Economic Review*, 39(1), 239–254.

Ebrahimy, E., and R. Shimer (2010): "Stock-flow matching," *Journal of Economic Theory*, 145(4), 1325–1353.

Faberman, R. J., and M. Kudlyak (forthcoming): "The Intensity of Job Search and Search Duration," *American Economic Journal: Macroeconomics*.

Gee, L. K. (2015): "Information Effects on Job Application Rates by Gender in a Large Field Experiment," Mimeo, Tufts University.

Gomme, P., and D. Lkhagvasuren (2015): "Worker search effort as an amplification mechanism," *Journal of Monetary Economics*, 75, 106–122.

Gregg, P., and B. Petrongolo (2005): "Stock-flow matching and the performance of the labor market," *European Economic Review*, 49(8), 1987–2011.

Herz, B., and T. van Rens (2015): "Accounting for Mismatch Unemployment," Discussion paper, University of Warwick.

Hornstein, A., P. Krusell, and G. L. Violante (2011): "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," *The American Economic Review*, 101(7).

Jaccard, P. (1901): "Étude de la distribution florale dans une portion des Alpes et du Jura," *Bulletin de la Societe Vaudoise des Sciences Naturelles*, 37, 547–579.

Jolivet, G., B. Jullien, and F. Postel-Vinay (2016): "Reputation and Prices on the e-Market: Evidence from a Major French Platform," *International Journal of Industrial Organization, forthcoming*.

Jolivet, G., and H. Turon (2014): "Consumer Search Costs and Preferences on the Internet," Mimeo, University of Bristol.

Kudlyak, M., D. Lkhagvasuren, and R. Sysuyev (2013): "Systematic Job Search: New Evidence from Individual Job Application Data," mimeo, Federal Reserve Bank of Richmond.

Kuhn, P., and H. Mansour (2014): "Is Internet Job Search still Ineffective?," *The Economic Journal*, 124(581), 1213–1233.

Lewis, G. (2011): "Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors," *The American Economic Review*, 101(4), 1535–1546.

Leyva, G. (2018): "Against All Odds: Job Search During the Great Recession," Unpublished Manuscript. Banco de Mexico.

Manning, A., and B. Petrongolo (2017): "How local are labor markets? Evidence from a spatial job search model," *American Economic Review*, 107(10), 2877–2907.

Marinescu, I., and R. Rathelot (2018): "Mismatch unemployment and the geography of job search," *American Economic Journal: Macroeconomics*, 10(3), 42–70.

Marinescu, I. E., and R. Rathelot (2015): "Mismatch Unemployment and the Geography of Job Search," Mimeo.

Marinescu, I. E., and R. P. Wolthoff (2015): "Opening the Black Box of the Matching Function: The Power of Words," Discussion Paper 9071, IZA.

Menzio, G., I. A. Telyukova, and L. Visschers (2016): "Directed search over the life cycle," *Review of Economic Dynamics*, 19, 38 – 62, Special Issue in Honor of Dale Mortensen.

Michelacci, C., and H. Ruffo (2015): "Optimal Life Cycle Unemployment Insurance," *American Economic Review*, 105(2), 816–59.

Michelacci, C., and J. Suarez (2006): "Incomplete Wage Posting," *Journal of Political Economy*, 114(6), 1098–1123.

Mukoyama, T., C. Patterson, and A. Şahin (2018): "Job Search Behavior over the Business Cycle," *American Economic Journal: Macroeconomics*, 10(1), 190–215.

Naudon, A., and A. Pérez (2018): "Unemployment dynamics in Chile: 1960-2015," *Journal Economía Chilena (The Chilean Economy)*, 21(1), 4–33.

Şahin, A., J. Song, G. Topa, and G. L. Violante (2014): "Mismatch unemployment," *American Economic Review*, 104(11), 3529–64.

Taylor, C. R. (1995): "The Long Side of the Market and the Short End of the Stick: Bargaining Power and Price Formation in Buyers', Sellers', and Balanced Markets," *The Quarterly Journal of Economics*, 110(3), 837–855.

Van Hooft, E. A., M. P. Born, T. W. Taris, and H. van der Flier (2005): "Predictors and outcomes of job search behavior: The moderating effects of gender and family situation," *Journal of Vocational Behavior*, 67(2), 133–152.

Van Nierop, E., B. Bronnenberg, R. Paap, M. Wedel, and P. H. Franses (2010): "Retrieving unobserved consideration sets from household panel data," *Journal of Marketing Research*, 47(1), 63–74.

van Ommeren, J., and G. Russo (2013): "Firm Recruitment Behaviour: Sequential or Non-sequential Search?," *Oxford Bulletin of Economics and Statistics*, pp. 1–24.

van Ours, J., and G. Ridder (1992): "Vacancies and the Recruitment of New Employees," *Journal of Labor Economics*, 10(2), 138–155.

Wiswall, M., and B. Zafar (2018): "Preference for the Workplace, Investment in Human Capital, and Gender," *The Quarterly Journal of Economics*, 133(1), 457–507.

Robustness for different subsamples

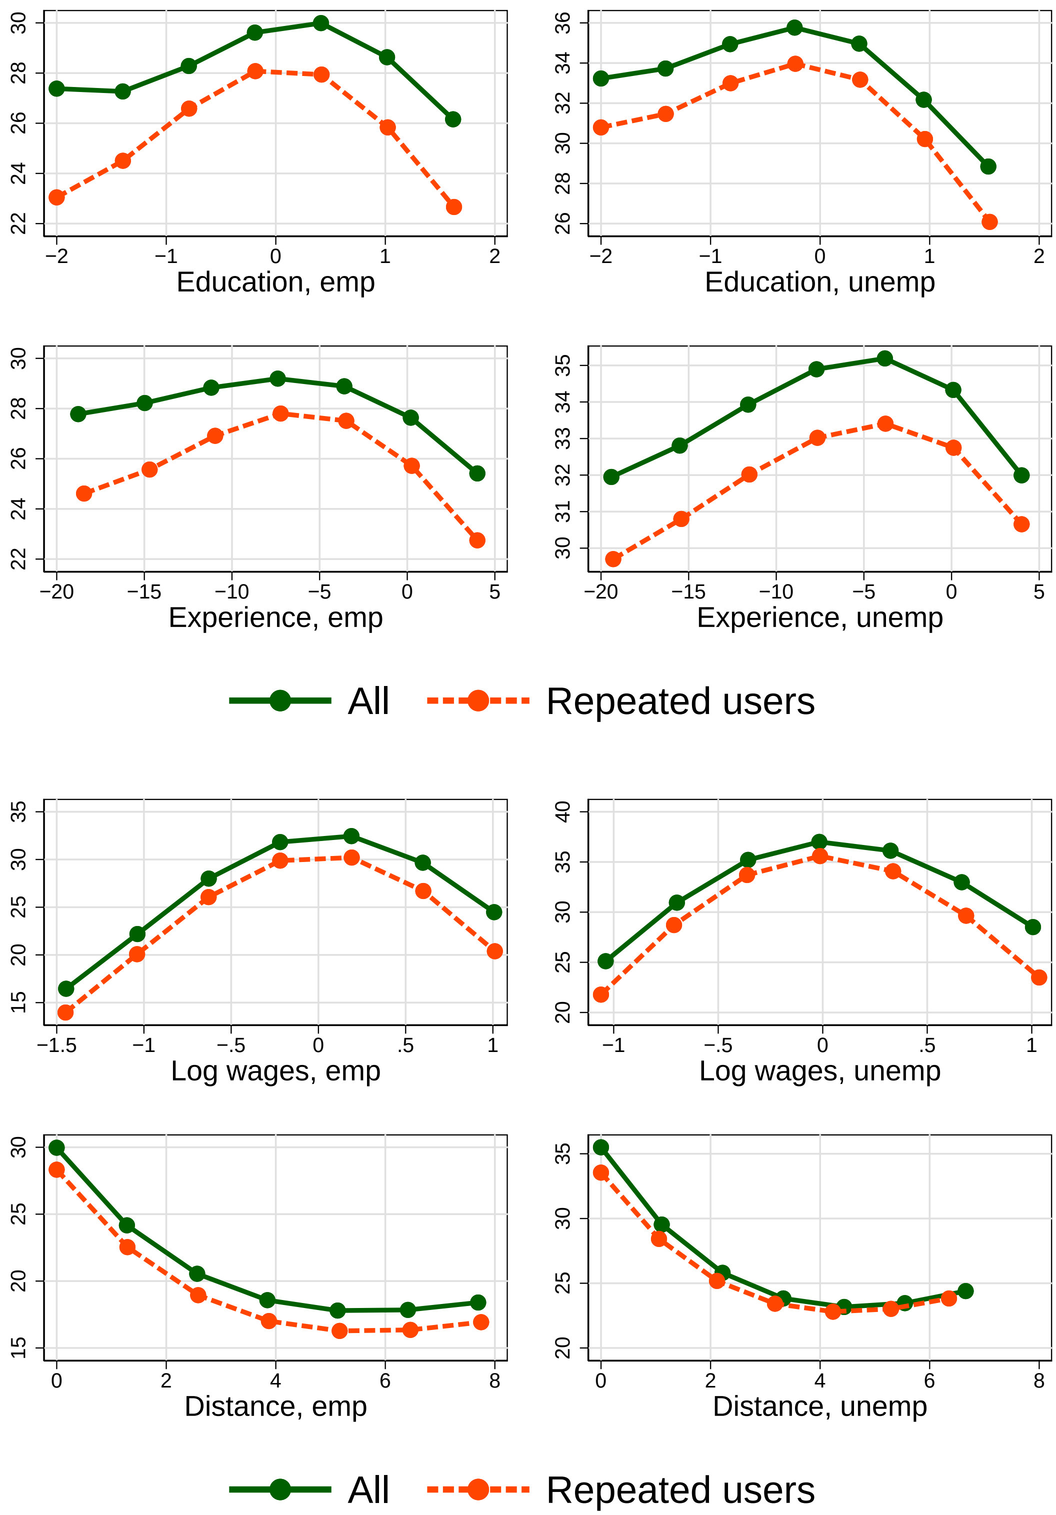

Figure A1: Contrasting predicted application probabilities for the whole sample and for repeated users of the website, given results from equation (3) and different levels of misalignment in selected variables: education, experience, log wages, and distance (see main text for details). The rest of regressors are at their sample means.

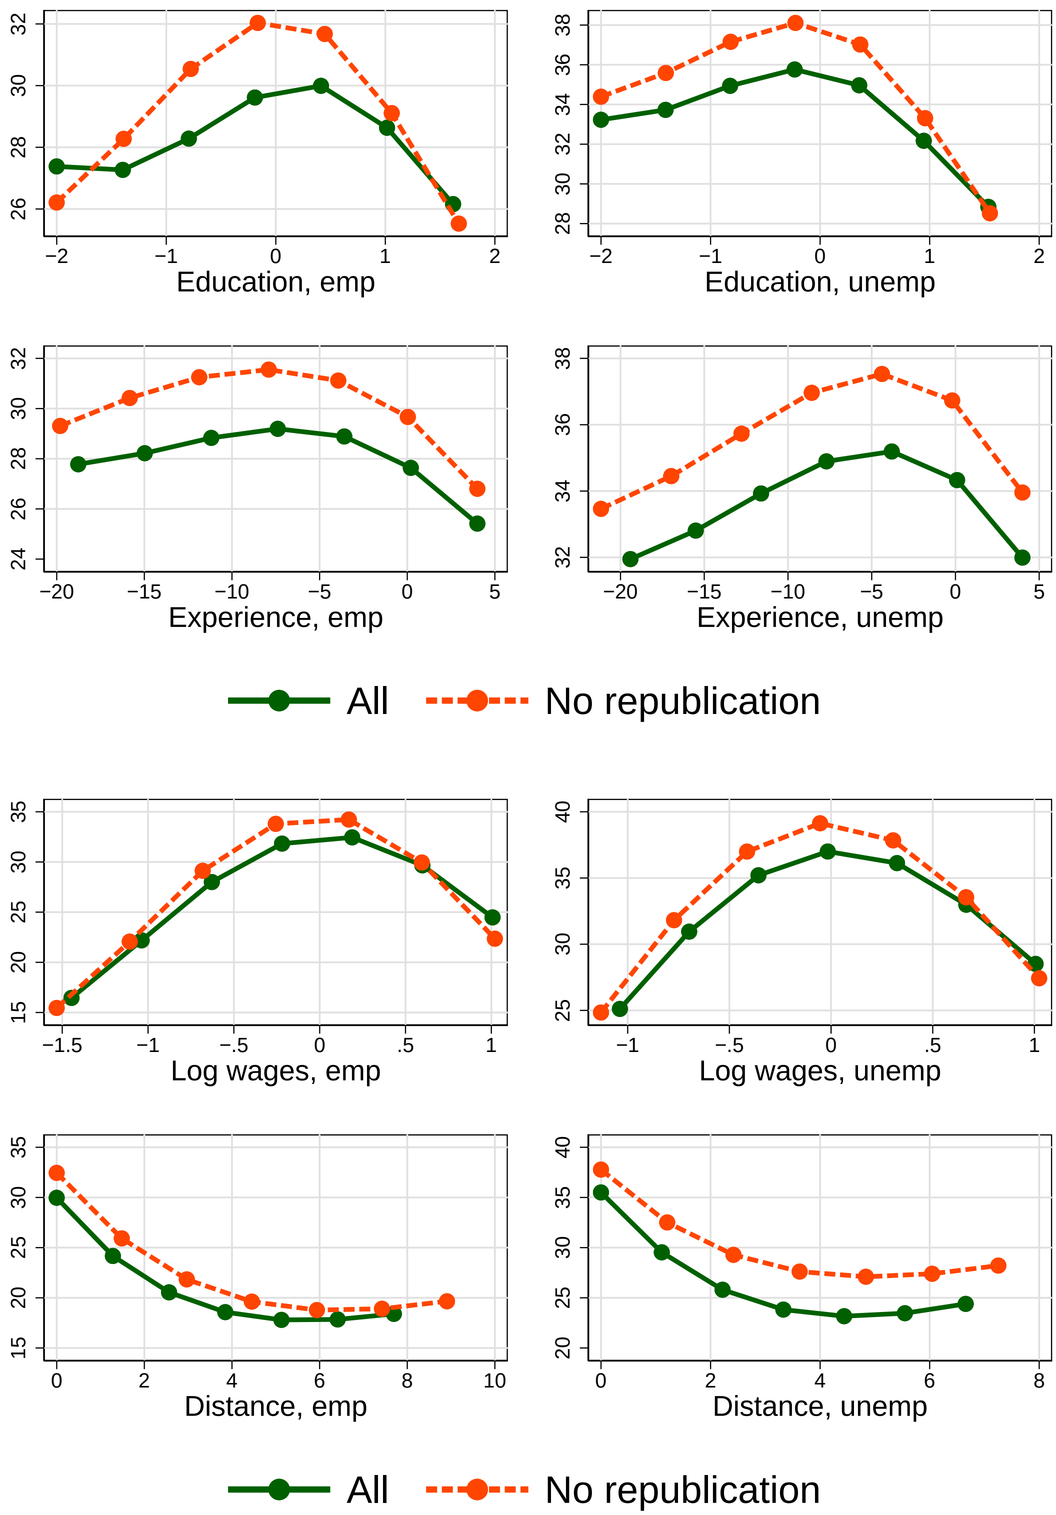

Figure A2: Contrasting predicted application probabilities for the whole sample and those for ads with no republication, given results from equation (3) and different levels of misalignment in selected variables: education, experience, log wages, and distance (see main text for details). The rest of regressors are at their sample means.

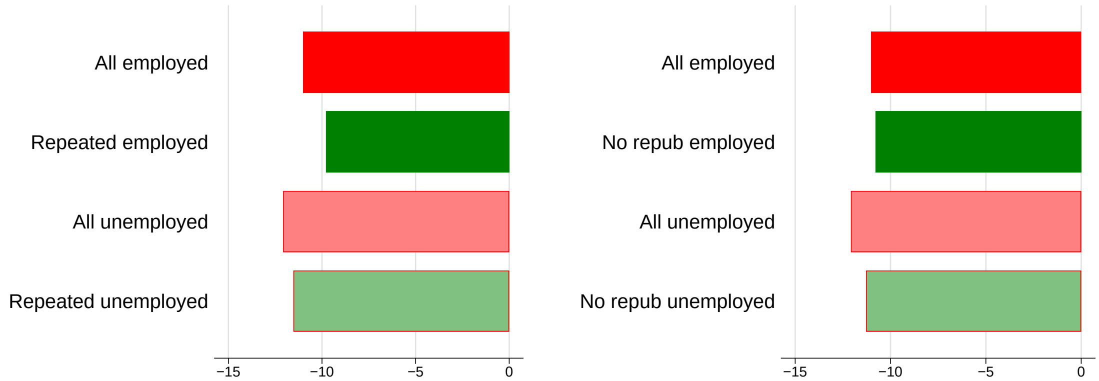

Figure A3: Contrasting predicted effects on application probabilities for the whole sample and for repeated users of the website, given results from equation (3) and different levels of misalignment in selected variables: education, experience, log wages, and distance (see main text for details). The rest of regressors are at their sample means.

Measures of probability with respect to average probability in reference group

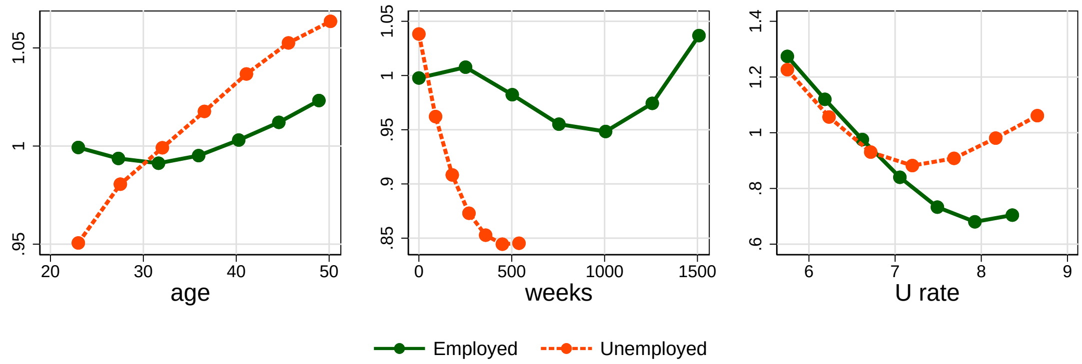

Figure A4: Predicted application probabilities (relative to sample averages) for different ages, number of weeks in the current labor force status and national unemployment rate at the time of the application decision, given results from eq. (3). The figure is computed using the coefficients associated to a polynomial of order 5 on each variable and leaving the rest of regressors at their sample mean.

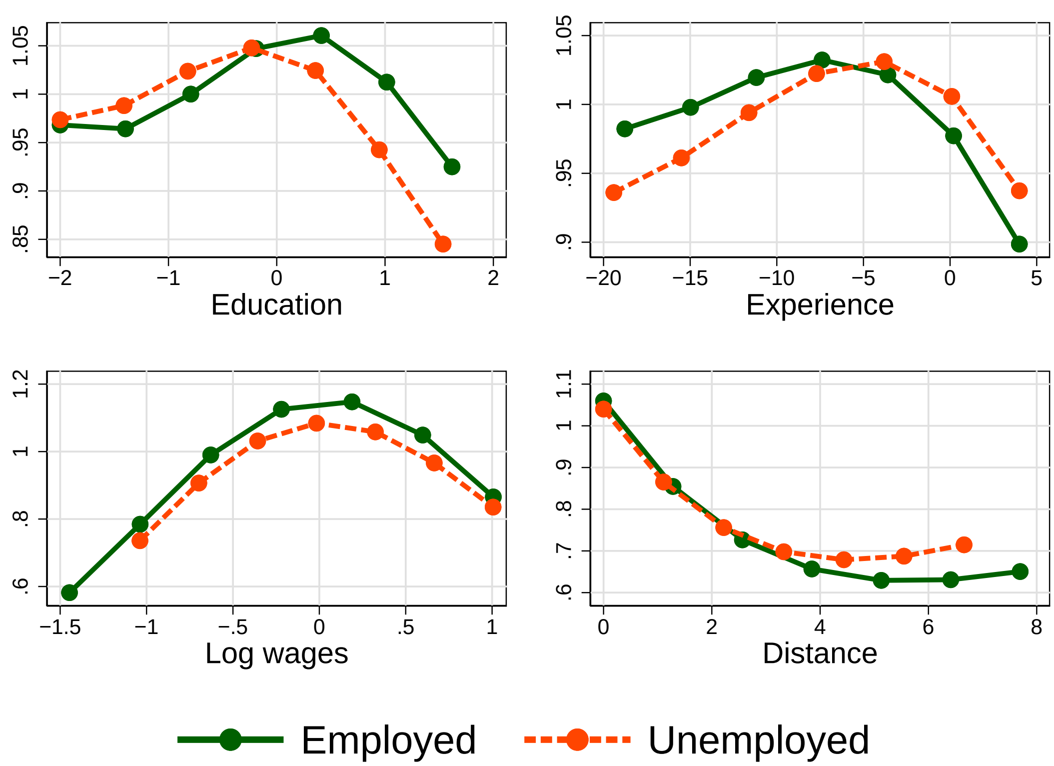

Figure A5: Predicted application probabilities, given results from eq. (3) and different levels of misalignment in the selected variable x (see main text for details). The rest of regressors are at their sample means. Results are relative to average application probability in the the corresponding group.

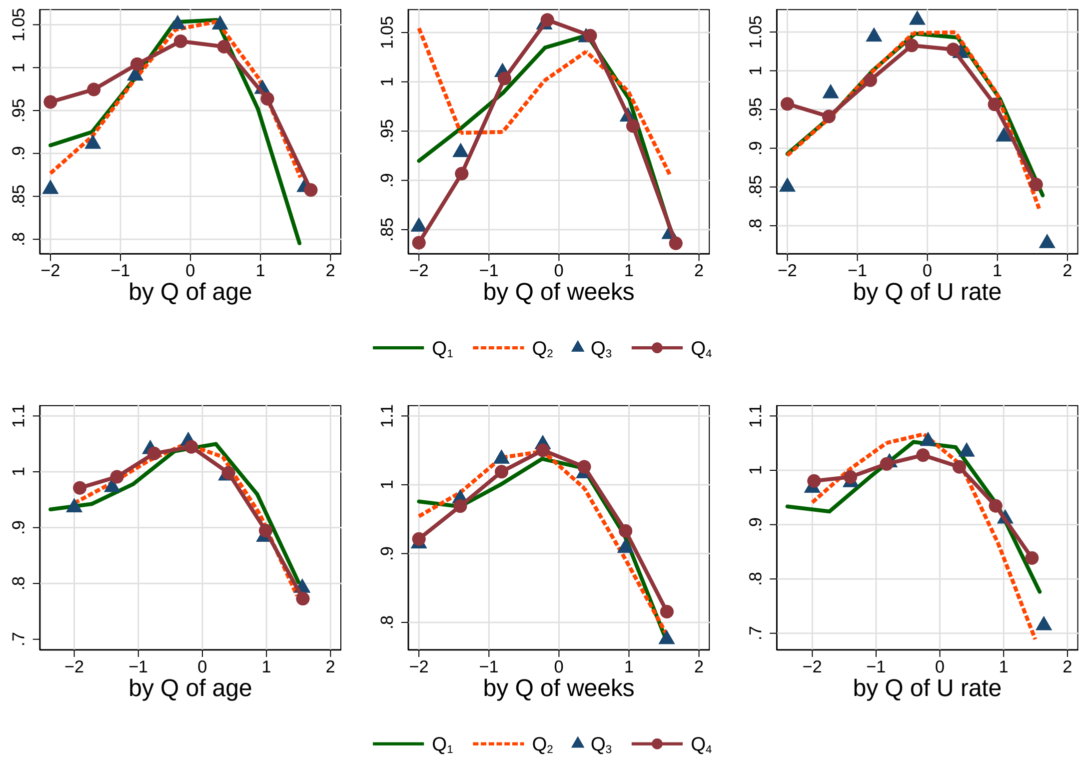

Figure A6: Predicted application probabilities for EMPLOYED (first row) and UNEMPLOYED (second row), given results from equation (3) and different levels of misalignment in EDUCATION (see main text for details). The rest of regressors are at their sample means. **Results are relative to average application probability in the the corresponding group.** The figure shows results when we split the sample in quartiles of age, weeks in the current labor force status, and the national unemployment rate.

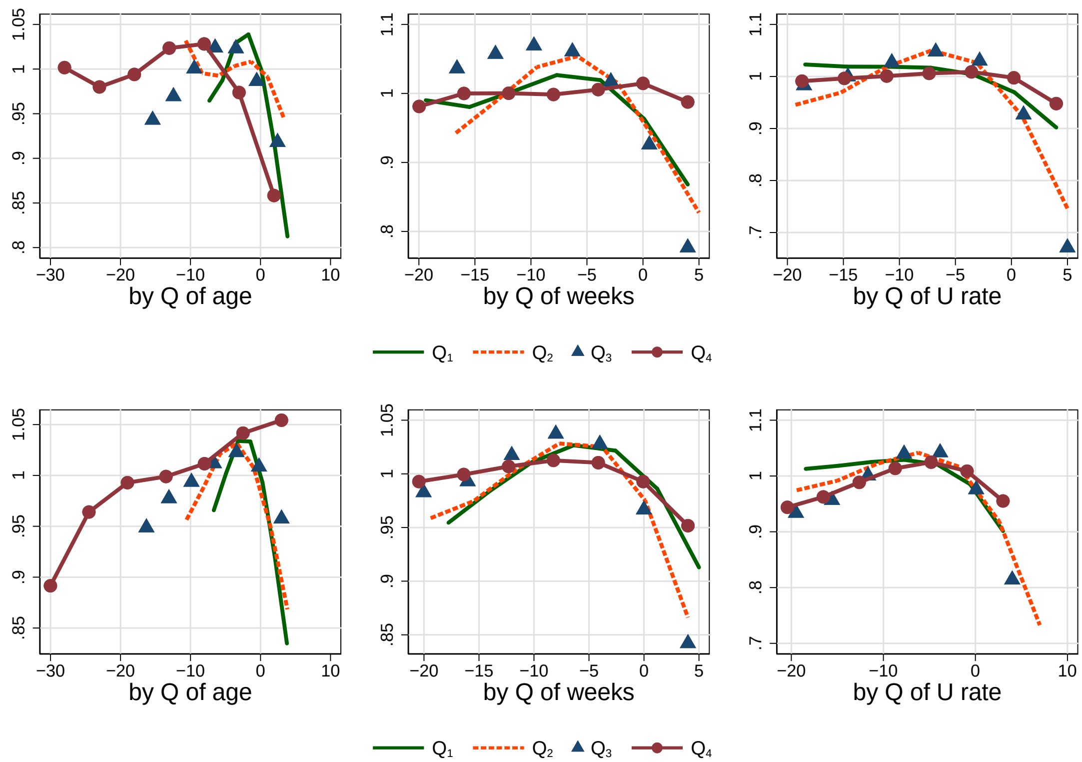

Figure A7: Predicted application probabilities for EMPLOYED (first row) and UNEMPLOYED (second row), given results from eq. (3) and different levels of misalignment in EXPERIENCE (see main text for details). The rest of regressors are at their sample means. **Results are relative to average application probability in the the corresponding group.** The figure shows results when we split the sample in quartiles of age, weeks in the current labor force status and national unemployment rate.

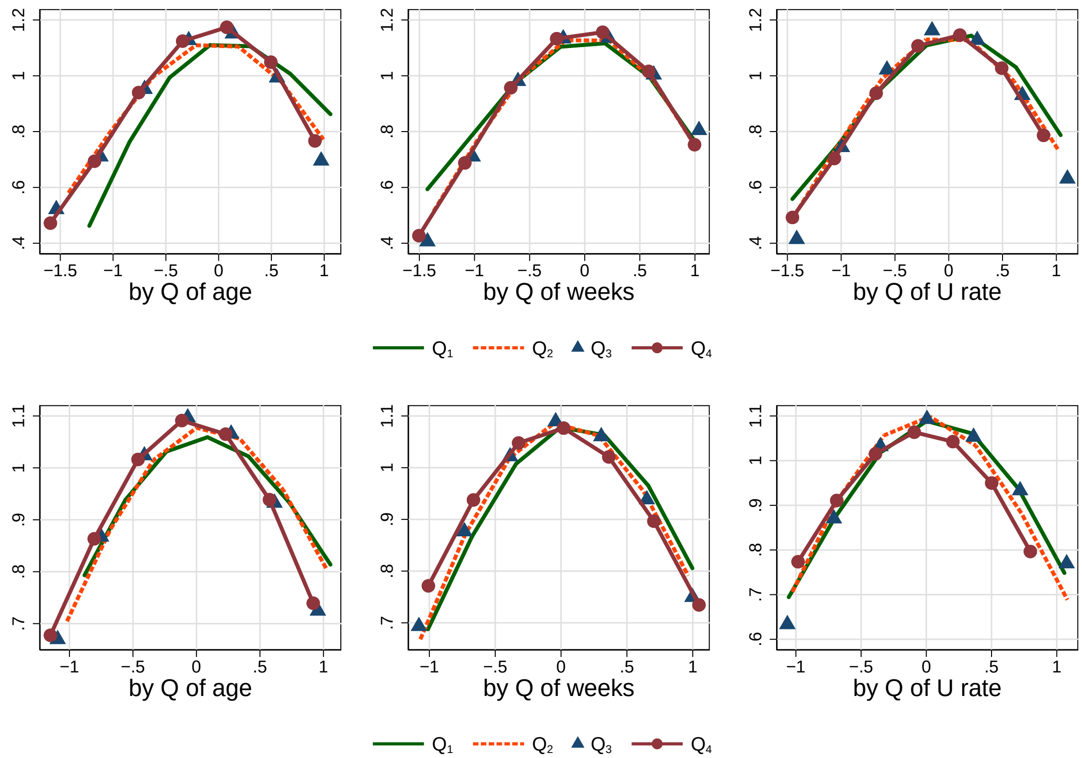

Figure A8: Predicted application probabilities for EMPLOYED (first row) and UNEMPLOYED (second row), given results from eq. (3) and different levels of misalignment in LOG-WAGES (see main text for details). The rest of regressors are at their sample means. **Results are relative to average application probability in the the corresponding group.** The figure shows results when we split the sample in quartiles of age, weeks in the current labor force status and national unemployment rate.

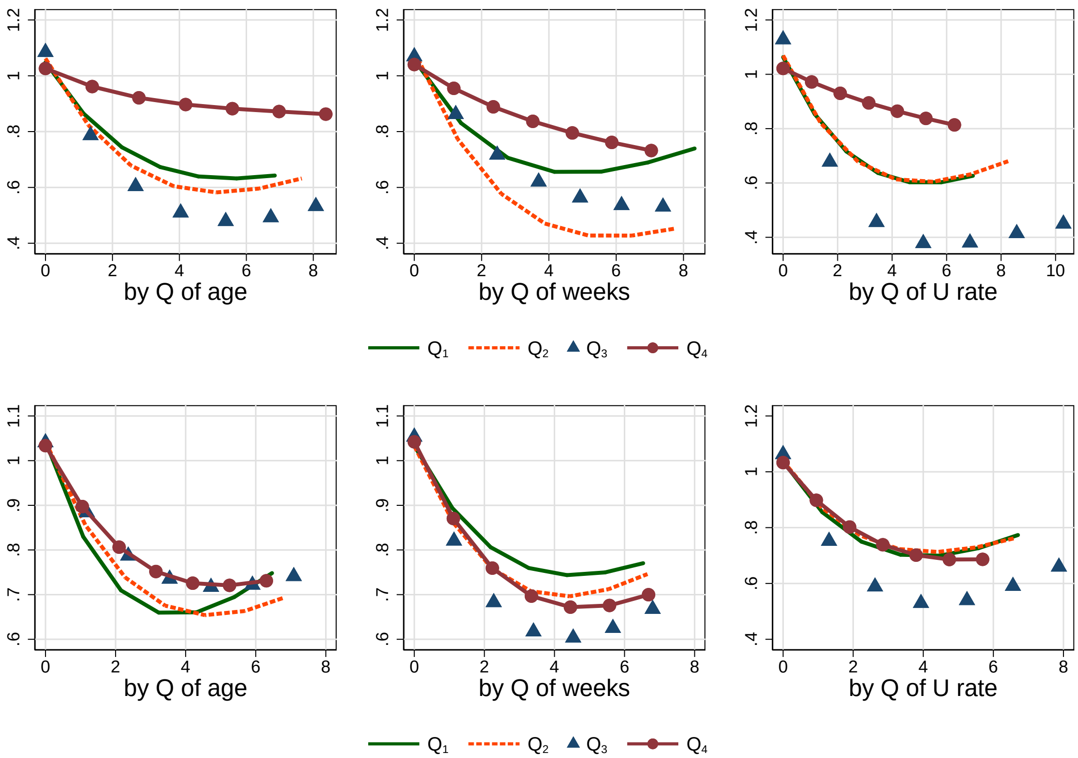

Figure A9: Predicted application probabilities for EMPLOYED (first row) and UNEMPLOYED (second row), given results from equation (3) and different levels of misalignment in DISTANCE (see main text for details). The rest of regressors are at their sample means. **Results are relative to average application probability in the the corresponding group.** The figure shows results when we split the sample in quartiles of age, weeks in the current labor force status, and the national unemployment rate.

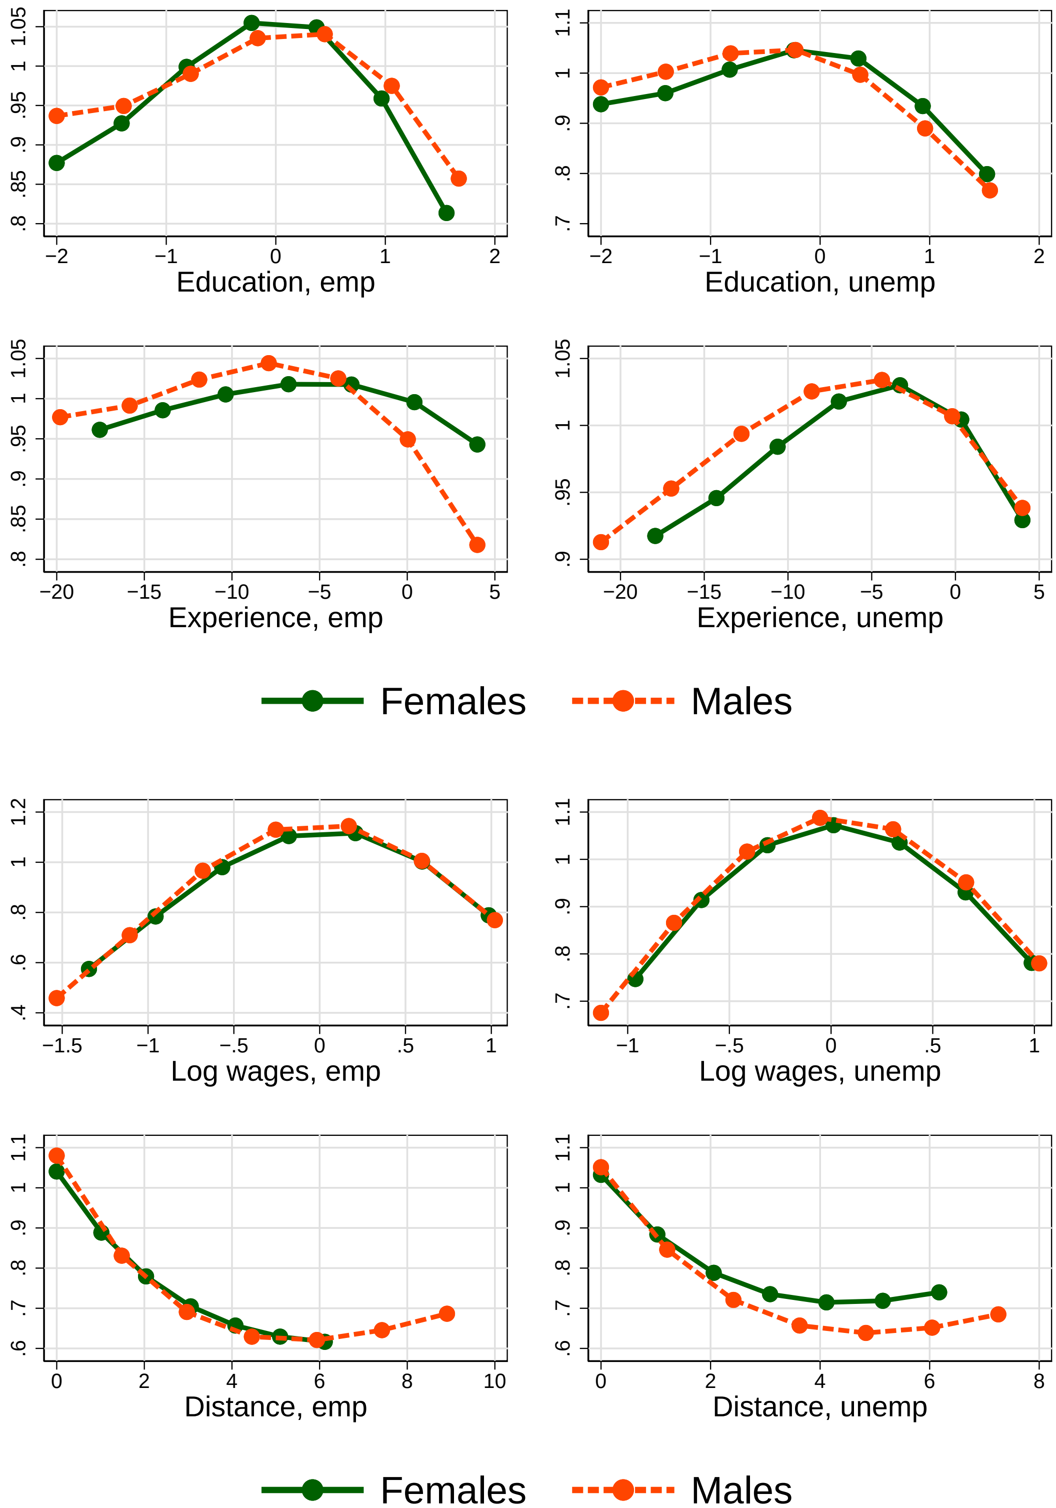

Figure A10: Predicted application probabilities, given results from equation 3 by gender and different levels of misalignment in the selected variable in selected variables: education, experience, log wages, and distance (see main text for details). The rest of regressors are at their sample means. **Results are relative to average application probability in the the corresponding group.**

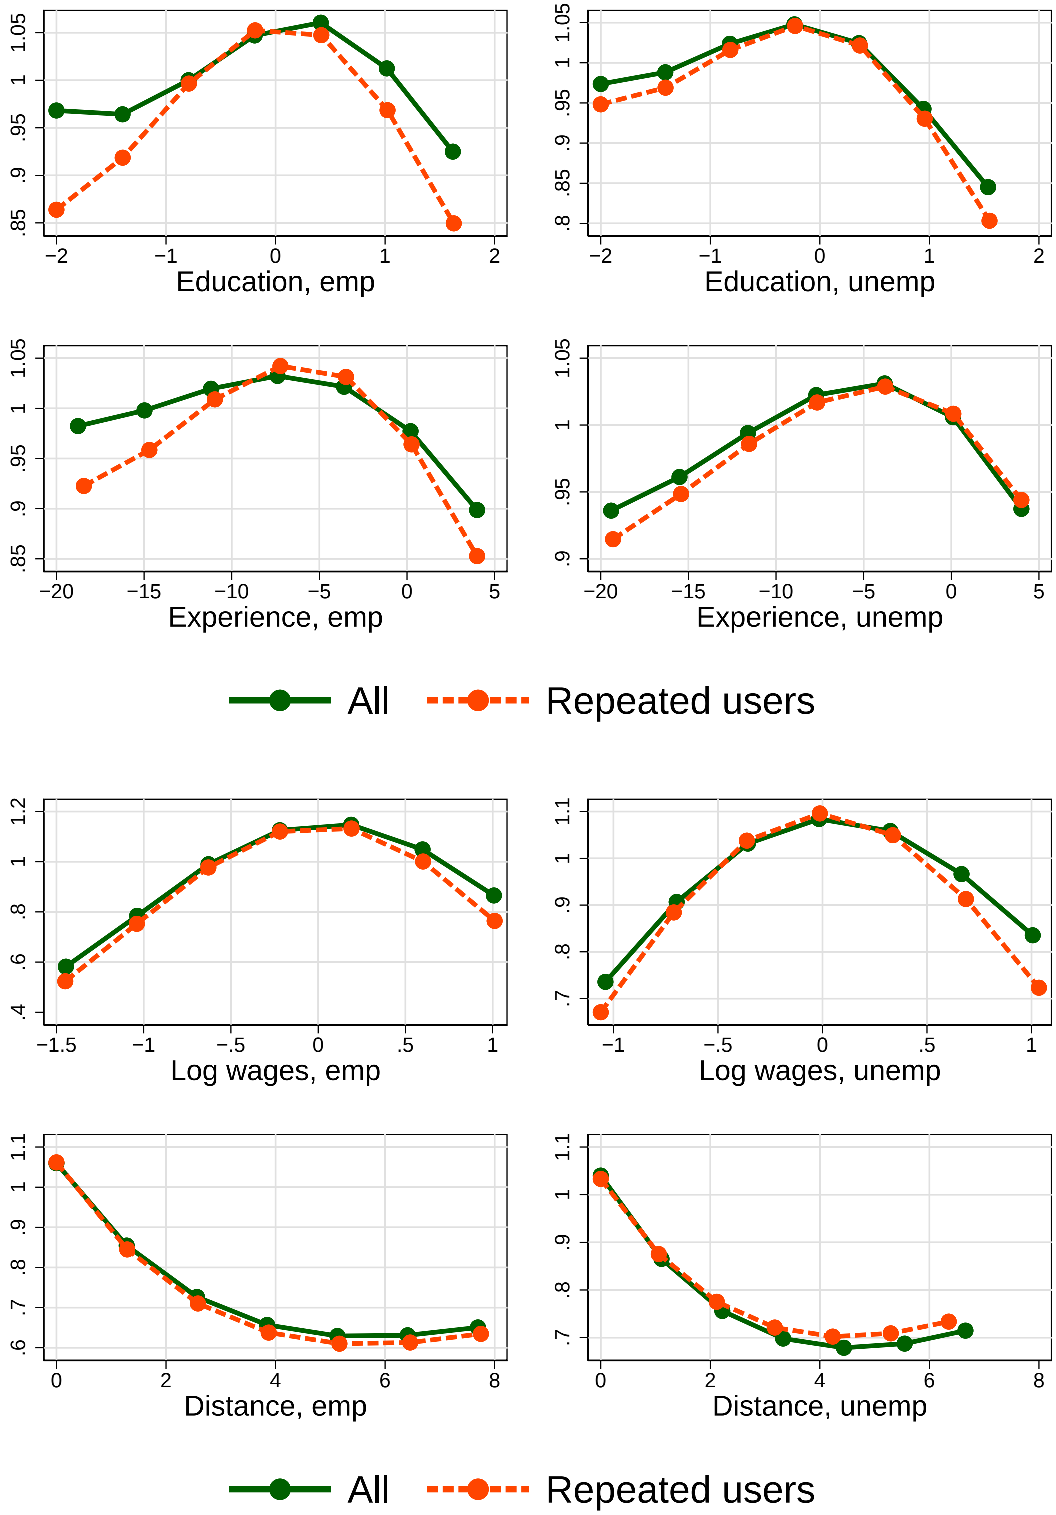

Figure A11: Contrasting predicted application probabilities for the whole sample and for repeated users of the website, given results from eq. (3) and different levels of misalignment in in selected variables: education, experience, log wages, and distance (see main text for details). The rest of regressors are at their sample means. **Results are relative to average application probability in the the corresponding group.**

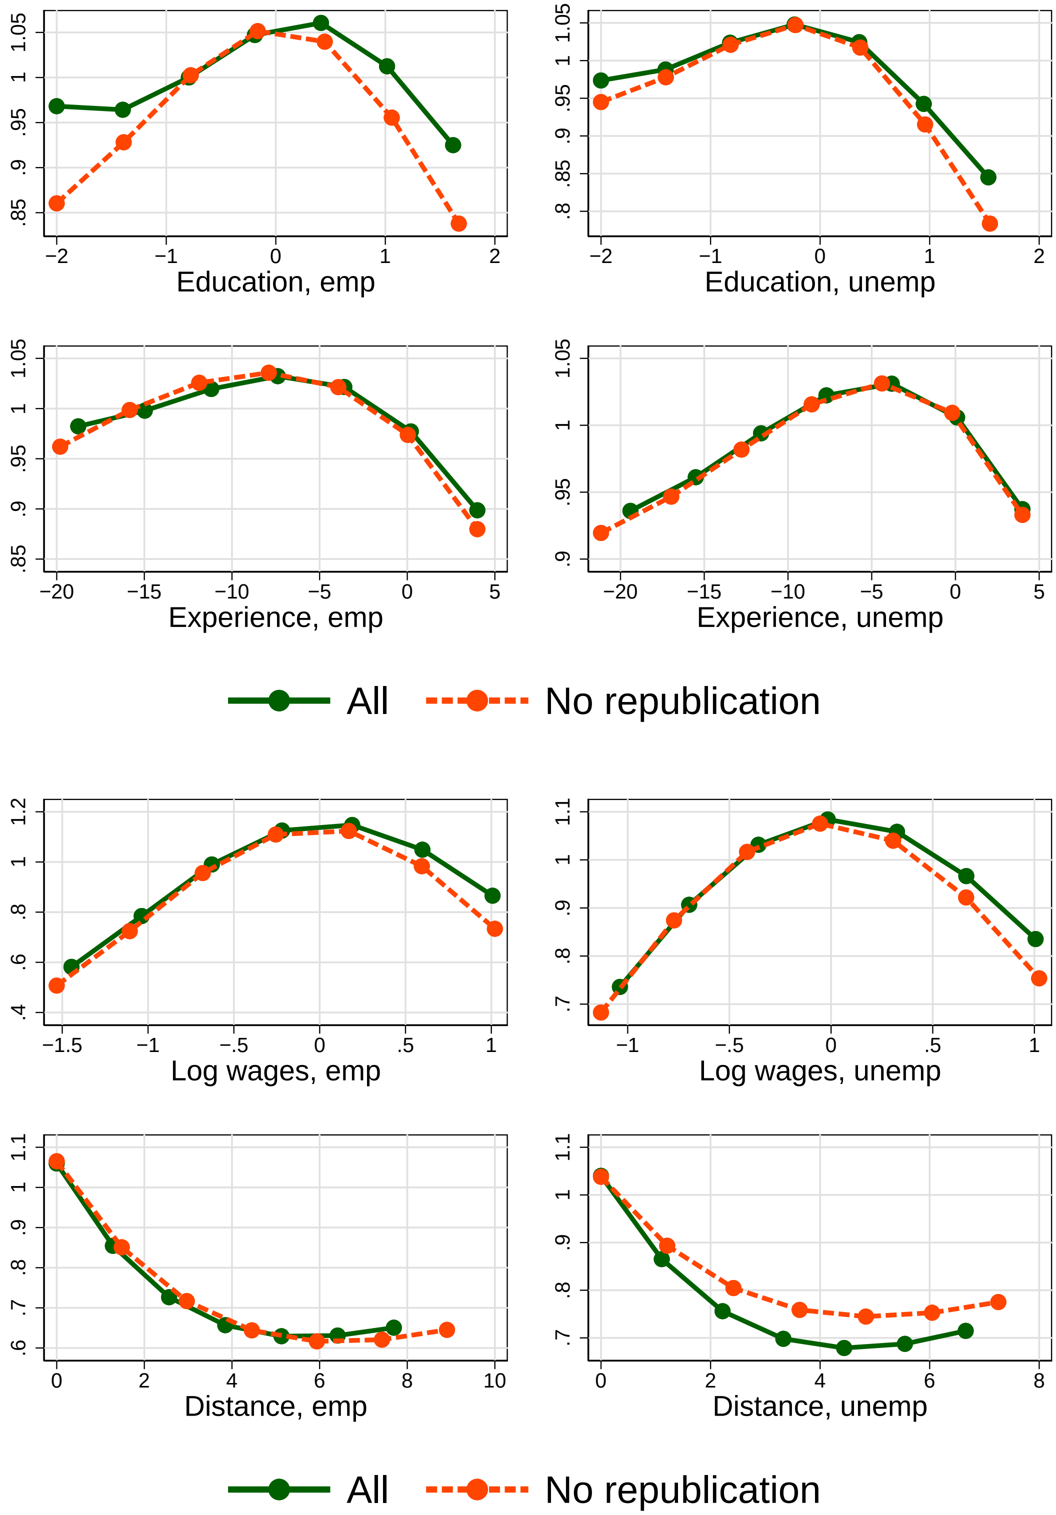

Figure A12: Contrasting predicted application probabilities for the whole sample and for the ads with no republication, given results from eq. (3) and different levels of misalignment in selected variables: education, experience, log wages, and distance (see main text for details). The rest of regressors are at their sample means. **Results are relative to average application probability in the the corresponding group.**

Table A1: Intensive margin coefficients: Robustness analysis

VARIABLES	(1) Employed No repub	(2) Employed Repeat	(3) Unemployed No repub	(4) Unemployed Repeat
Married	6.852***	8.062***	-1.731***	0.752*
	(0.678)	(0.663)	(0.434)	(0.408)
Male	0.926***	1.753***	1.169***	2.019***
	(0.087)	(0.070)	(0.074)	(0.063)
Male x Married	-1.179***	-1.387***	1.625***	0.958***
	(0.143)	(0.119)	(0.134)	(0.117)
Explicit wage (w)	0.245***	0.098*	-0.420***	0.223***
	(0.068)	(0.056)	(0.060)	(0.052)
Explicit wage (a)	-2.319***	-2.962***	0.569***	0.369***
	(0.109)	(0.090)	(0.081)	(0.071)
No. of Vacancies (a)	-0.016***	-0.041***	0.022***	0.007***
	(0.001)	(0.001)	(0.001)	(0.001)
Ad duration (weeks)	-0.040***	0.003	-0.256***	-0.242***
	(0.006)	(0.003)	(0.005)	(0.003)
Observations	1,691,836	2,359,611	2,355,090	3,019,631
R-squared	0.135	0.138	0.136	0.128

Notes: Regression coefficients from a linear regression on application decisions. Dependent variable is y_{aw}, a dummy for the existence of a job application. Each regression controls also for polynomials and interactions in *misalignment* as well as age of the worker, firm size, contract type, dummies for different types of requirements of the job and characteristics of the firm (see details in the main text). Standard errors in parentheses. One, two, and three asterisks indicate significance at 10%,5%, and 1%, respectively.

Table A2: Intensive margin coefficients by age quartile and labor status

VARIABLES	(1) Employed Q=1	(2) Employed Q=2	(3) Employed Q=3	(4) Employed Q=4	(5) Unemployed Q=1	(6) Unemployed Q=2	(7) Unemployed Q=3	(8) Unemployed Q=4
Married	-22.619***	-15.212	14.252***	3.435***	-16.913**	-12.243***	9.824***	1.486***
	(6.604)	(10.227)	(1.532)	(0.508)	(8.037)	(2.210)	(0.865)	(0.340)
Male	2.614***	0.642***	1.736***	-0.947***	1.149***	1.051***	3.175***	1.916***
	(0.105)	(0.128)	(0.136)	(0.161)	(0.091)	(0.096)	(0.119)	(0.140)
Male x Married	-1.121***	-0.414*	-2.382***	0.301	4.387***	2.912***	-2.381***	0.943***
	(0.289)	(0.235)	(0.202)	(0.204)	(0.334)	(0.217)	(0.192)	(0.179)
Explicit wage (w)	-0.990***	0.250**	0.313***	0.803***	0.635***	-0.793***	-0.199**	-0.074
	(0.095)	(0.106)	(0.102)	(0.101)	(0.084)	(0.083)	(0.093)	(0.088)
Explicit wage (a)	-0.301**	-1.246***	-5.069***	-1.604***	2.268***	0.012	-0.951***	0.136
	(0.145)	(0.171)	(0.167)	(0.160)	(0.113)	(0.114)	(0.127)	(0.118)
No. of Vacancies (a)	0.024***	-0.037***	-0.047***	-0.034***	0.045***	0.020***	0.007***	0.029***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)	(0.002)	(0.002)
Ad duration (weeks)	-0.008	0.040***	0.066***	0.008	-0.122***	-0.225***	-0.254***	-0.219***
	(0.005)	(0.006)	(0.006)	(0.006)	(0.005)	(0.005)	(0.006)	(0.006)
Observations	794,285	671,103	735,774	754,214	1,092,849	1,162,565	992,684	1,082,161
R-squared	0.172	0.157	0.156	0.140	0.177	0.144	0.126	0.133

Notes: Regression coefficients from a linear regression on application decisions. Dependent variable is y_{aw}, a dummy for the existence of a job application. Each regression controls also for polynomials and interactions in *misalignment* as well as age of the worker, firm size, contract type, dummies for different types of requirements of the job and characteristics of the firm (see details in the main text). Standard errors in parentheses. One, two, and three asterisks indicate significance at 10%,5%, and 1%, respectively.

Table A3: Intensive margin coefficients by duration and labor status

VARIABLES	(1) Employed Q=1	(2) Employed Q=2	(3) Employed Q=3	(4) Employed Q=4	(5) Unemployed Q=1	(6) Unemployed Q=2	(7) Unemployed Q=3	(8) Unemployed Q=4
Married	13.208***	2.051	-9.538***	-0.703	0.409	-1.553**	8.710***	1.147**
	(0.830)	(1.787)	(0.996)	(0.881)	(0.760)	(0.664)	(0.613)	(0.538)
Male	1.468***	1.524***	1.086***	1.515***	0.751***	2.385***	1.858***	1.085***
	(0.107)	(0.178)	(0.114)	(0.128)	(0.106)	(0.106)	(0.103)	(0.109)
Male x Married	-2.479***	-2.421***	0.920***	-0.920***	0.903***	-1.189***	1.975***	2.240***
	(0.177)	(0.338)	(0.209)	(0.200)	(0.200)	(0.201)	(0.193)	(0.194)
Explicit wage (w)	0.428***	-1.937***	-0.102	-0.310***	0.657***	-0.847***	-0.281***	0.160*
	(0.084)	(0.149)	(0.094)	(0.099)	(0.087)	(0.088)	(0.085)	(0.086)
Explicit wage (a)	-2.422***	-1.585***	-2.559***	-1.306***	0.997***	-0.341***	0.939***	0.818***
	(0.135)	(0.232)	(0.153)	(0.155)	(0.118)	(0.120)	(0.116)	(0.116)
No. of Vacancies (a)	-0.005***	-0.024***	-0.025***	-0.025***	0.041***	0.009***	0.004***	0.031***
	(0.002)	(0.003)	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)	(0.001)
Ad duration (weeks)	-1.012***	0.189***	0.187***	0.133***	-0.043***	-0.243***	-0.280***	-0.218***
	(0.009)	(0.008)	(0.005)	(0.005)	(0.006)	(0.006)	(0.005)	(0.005)
Observations	1,193,917	288,595	751,106	721,758	1,111,480	1,060,613	1,099,879	1,058,287
R-squared	0.159	0.159	0.130	0.122	0.141	0.138	0.144	0.144

Notes: Regression coefficients from a linear regression on application decisions. Dependent variable is y_{aw}, a dummy for the existence of a job application. Each regression controls also for polynomials and interactions in *misalignment* as well as age of the worker, firm size, contract type, dummies for different types of requirements of the job and characteristics of the firm (see details in the main text). Standard errors in parentheses. One, two, and three asterisks indicate significance at 10%,5%, and 1%, respectively.

Table A4: Intensive margin coefficients by unemployment rate quartile and labor status

VARIABLES	(1) Employed Q=1	(2) Employed Q=2	(3) Employed Q=3	(4) Employed Q=4	(5) Unemployed Q=1	(6) Unemployed Q=2	(7) Unemployed Q=3	(8) Unemployed Q=4
Married	8.119***	-1.690**	7.893***	6.767***	2.329***	4.283***	-2.339***	1.546*
	(0.851)	(0.830)	(1.195)	(1.315)	(0.600)	(0.568)	(0.544)	(0.858)
Male	-0.727***	1.999***	1.117***	4.292***	0.390***	1.416***	2.481***	2.778***
	(0.118)	(0.103)	(0.124)	(0.170)	(0.097)	(0.101)	(0.086)	(0.149)
Male x Married	-0.140	-3.015***	0.863***	-1.474***	1.471***	-0.361*	3.255***	1.106***
	(0.201)	(0.172)	(0.217)	(0.275)	(0.188)	(0.187)	(0.162)	(0.258)
Explicit wage (w)	-0.267***	0.135*	0.135	-0.134	-0.717***	-0.211**	1.480***	0.479***
	(0.094)	(0.081)	(0.098)	(0.135)	(0.080)	(0.083)	(0.071)	(0.117)
Explicit wage (a)	-1.216***	-1.694***	-2.033***	-2.698***	1.628***	-0.054	-0.717***	-0.365**
	(0.159)	(0.133)	(0.164)	(0.232)	(0.113)	(0.118)	(0.105)	(0.182)
No. of Vacancies (a)	0.030***	-0.030***	-0.108***	-0.044***	0.009***	0.032***	0.071***	0.061***
	(0.002)	(0.002)	(0.007)	(0.003)	(0.001)	(0.002)	(0.002)	(0.002)
Ad duration (weeks)	0.048***	-0.017***	0.015**	-0.061***	-0.264***	-0.276***	-0.211***	-0.132***
	(0.006)	(0.005)	(0.006)	(0.009)	(0.006)	(0.005)	(0.004)	(0.008)
Observations	831,942	1,078,839	557,788	486,807	1,234,245	1,093,609	1,312,733	689,672
R-squared	0.136	0.147	0.099	0.168	0.146	0.121	0.181	0.101

Notes: Regression coefficients from a linear regression on application decisions. Dependent variable is y_{aw}, a dummy for the existence of a job application. Each regression controls also for polynomials and interactions in *misalignment* as well as age of the worker, firm size, contract type, dummies for different types of requirements of the job and characteristics of the firm (see details in the main text). Standard errors in parentheses. One, two, and three asterisks indicate significance at 10%,5%, and 1%, respectively.